

PCI Express External Cabling Specification

Revision 3.0, Version 0.9

April 12, 2019



PCI-SIG® disclaims all warranties and liability for the use of this document and the information contained herein and assumes no responsibility for any errors that appear in this document, nor does PCI-SIG make a commitment to update the information contained herein.

Contact the PCI-SIG office to obtain the latest revision of this specification.

Forward any questions regarding the PCI Code and ID Assignment Specification or membership in PCI-SIG to:

Membership Services

www.pcisig.com

E-mail: administration@pcisig.com

Phone: 503-619-0569

Fax: 503-644-6708

Technical Support

techsupp@pcisig.com

DISCLAIMER

This *PCI Express External Cabling Specification* is provided *as is* with no warranties whatsoever, including any warranty of merchantability, non-infringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification, or sample. PCI-SIG disclaims all liability for infringement of proprietary rights, relating to use of information in this specification. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted herein.

Trademarks

PCI, PCI Express, PCIe, ExpressModule, and PCI-SIG are trademarks or registered trademarks of PCI SIG. Other company and product names are trademarks of the respective companies with which they are associated.

Copyright

© 2007 – 2019 PCI SIG. All rights reserved.

Revision History

Revision	Version	History	Date
1.0		Initial public release	January 7, 2007
2.0		Release includes 5.0 GT/s	July 22, 2012
3.0	0.1	Initial working draft	June 19, 2013
	0.3	Updated revision, updated header	October 10, 2013
	0.5	Following additions/updates were done: <ul style="list-style-type: none"> Added MGTPWR Added port configuration options Added sideband function Added SAS interop (preliminary) Added notes to the memory map, the new style guide is currently not being referenced, Updated 2-wire interface Updated diagrams Changed SMBus to CMI Swapped pin assignments for lanes 0 & 1 Removed the cable clock from the Upstream system completely Moved CREFCLK information to Appendix for Legacy Adapter Cable Updated images Updated terms and punctuation for consistency 	March 3, 2015
	0.6	<ul style="list-style-type: none"> Included feedback from 0.5 feedback Updated some pinout tables with correct names Removed notes column from cable memory map tables Added power sequencing requirements Updated cable performance requirements, included recommended settings for RxEQ based on cable loss 	May 25, 2015
	0.7	Updated the following: <ul style="list-style-type: none"> 6.2.2.3.7, Bytes 108-109, Propagation Delay Added Figure 6-3, Bytes 108-109, Propagation Delay Added Table 6-6, Lower Bound, Upper Bound, Base, and Span Values Updated equations in Chapter 6, Cable Specification Updated language per PCI SIG Style Guide 	November 9, 2016
	08	<ul style="list-style-type: none"> Incorporated Cable_30_errata_table_20170504 Incorporated PCI Express External Cabling R3.0_v0.7_TS_Edits 20170526 to Tech Pubs. Added the following sections Per PCI Express External Cabling R3.0_v0.8_20170713_toTP <ul style="list-style-type: none"> Added section 4.1.3, <i>Static Data CRC</i> Added section 5.3, <i>Fixed Side Connector Labeling</i> Added section 6.4, <i>Cable Assembly Labeling</i> 	August 8, 2017
	08	Incorporated <ul style="list-style-type: none"> PCI_Express_External_Cabling_R3.0_v0.8_TS_20170808_feedback 180323 Cable 3.0 0.9 Curves 	April 9, 2018
	09	<ul style="list-style-type: none"> Incorporated PCI Express External Cabling R3.0_v0.8_WG_181010 Update Figure 2-1, Figure 2-2, Figure 2-3 Figure 6-4, Figure 6-5, and Table 6-11 Updated 5 GT/s and 2.5 GT/s Eye Diagrams, changed crosstalk frequency limits to ICN limit 	April 812 , 2019

Table of Contents

11	1. INTRODUCTION.....	9
12	1.1. Terms and Acronyms	9
13	1.2. Reference Documents.....	11
14	1.3. Documentation Conventions	12
15	1.4. Specification Contents.....	12
16	1.5. Objectives	13
17	1.6. Overview.....	13
18	2. AUXILIARY SIGNALS	16
19	2.1. PWR.....	17
20	2.2. MGTPWR	17
21	2.3. CBLPRSNT#.....	18
22	2.4. CADDR.....	19
23	2.5. CINT#.....	20
24	2.6. Cable Management Interface	21
25	2.6.1. Capacitive Load of High-power Cable Management Interface Lines	25
26	2.6.2. Minimum Current Sinking Requirements for Cable Management Interface Devices.....	25
27	2.6.3. Cable Management Interface <i>Back Powering</i> Considerations.....	26
28	2.7. Electrostatic Discharge	26
29	2.8. Auxiliary Signal Parametric Specifications	27
30	2.8.1. DC Specifications for Auxiliary Signals	27
31	2.8.2. AC Specifications for Auxiliary Signals.....	28
32	2.9. Power Requirements.....	29
33	3. PCI EXPRESS SIGNALS	30
34	3.1. Interconnect	31
35	3.1.1. Link Definition	31
36	3.2. Electrical Budget.....	32
37	3.2.1. AC-coupling	32
38	3.2.2. Jitter from Clock Sources	32
39	3.2.3. Crosstalk.....	33
40	3.2.4. Lane-to-Lane Skew	33
41	3.2.5. Transmitter De-Emphasis.....	33
42	3.2.6. Skew Within the Differential Pair (Intra-Pair Skew).....	34
43	3.3. Eye Diagrams.....	34
44	3.3.1. Fixed-Port Transmit Compliance Eye.....	34
45	3.3.2. Cable Port Transmit Compliance Eye	36
46	3.4. Receiver Path Sensitivity Testing	38
47	3.5. ESD	39

48	4. INTEROPERABILITY	40
49	4.1. Cable Management Interface	41
50	4.1.1. Fixed-Side Memory Map	41
51	4.1.2. Sideband Messages	43
52	4.1.2.1. CMI_RESET	44
53	4.1.2.2. CMI_POWER_ENABLE	44
54	4.1.2.3. CMI_SSTART	45
55	4.1.2.4. CMI_CLP_READY	45
56	4.1.2.5. CMI_WAKE	45
57	4.1.2.6. CMI_CLP_PRESENT	45
58	4.1.2.7. CMI_HOT-PLUG_ATTENTION_BUTTON	45
59	4.1.2.8. CMI_HOT-PLUG_INDICATOR_REQ	46
60	4.1.2.9. Sideband Message CRC	46
61	4.1.2.10. Cable Port Status Register	46
62	4.1.2.11. Cable Port Status Register CRC	46
63	4.1.2.12. PCI Express Cable Port Capabilities Register	46
64	4.1.2.13. Static Data CRC	47
65	4.2. Power Sequencing	47
66	4.2.1. Power-Up Sequencing	47
67	4.2.1.1. Equalizer Settings Based on Cable Memory	53
68	4.2.1.2. PCI Express Configurations Based on Cable Memory	54
69	4.2.2. Power-Down Sequencing	55
70	4.2.2.1. Sequence Steps	55
71	4.2.3. Cable Management Interface Timings	55
72	4.2.4. Link Power Management	55
73	4.3. PCI Express Features	56
74	5. FIXED-SIDE DEFINITION	57
75	5.1. Fixed-Side Channel Characteristics	57
76	5.2. Signal Description	57
77	5.3. Fixed-Side Connector	58
78	5.3.1. Pin-out	58
79	5.4. Fixed Side Connector Labeling	61
80	6. CABLE SPECIFICATION	62
81	6.1. SFF-8644 Cable Assembly	62
82	6.2. Cable Management Interface	65
83	6.2.1. Memory Map	65
84	6.2.2. Memory Map Architecture	65
85	6.2.2.1. Bytes 0-127 (Lower Memory)	67
86	6.2.2.2. Bytes 128-255 ("Page 0, Upper Memory")	68
87	6.2.2.3. Field Descriptions	70
88	6.2.2.3.1. Bytes 1-2, Status Indicators	70
89	6.2.2.3.2. Byte 3, Interrupt Flags – LOS	70
90	6.2.2.3.3. Byte 6, Interrupt Flags – Temp Alarm	70
91	6.2.2.3.4. Byte 7, Interrupt Flags – Voltage Alarm	71
92	6.2.2.3.5. Bytes 22-23, Module Monitors – Temperature	71
93	6.2.2.3.6. Bytes 26-27, Module Monitors – Supply Voltage	71
94	6.2.2.3.7. Bytes 108-109, Propagation Delay	71
95	6.2.2.3.8. Byte 111, Cabled PCI Express Capabilities 1	73
96	6.2.2.3.9. Byte 112, Cabled PCI Express Capabilities 2	73

97	6.2.2.3.10. Bytes 119-126 Password Entry and Change	73
98	6.2.2.3.11. Byte 127 Page Select.....	74
99	6.2.2.3.12. Byte 147, Cable Technology.....	74
100	6.2.2.3.13. Bytes 148-163, Vendor Name.....	75
101	6.2.2.3.14. Bytes 165-166, PCI-SIG Vendor ID	75
102	6.2.2.3.15. Bytes 168-183, Vendor Part Number.....	75
103	6.2.2.3.16. Bytes 184-185, Vendor Revision	75
104	6.2.2.3.17. Bytes 186-189, Attenuation at Frequency.....	75
105	6.2.2.3.18. Byte 190, Maximum Case Temperature	76
106	6.2.2.3.19. Byte 191, Checksum-Base	76
107	6.2.2.3.20. Byte 195, Options.....	76
108	6.2.2.3.21. Bytes 196-211, Vendor Serial Number	76
109	6.2.2.3.22. Bytes 212-217, Vendor Date Code	76
110	6.2.2.3.23. Bytes 218-219, Vendor Lot Code.....	76
111	6.2.2.3.24. Byte 223, Checksum-Extended	76
112	6.2.2.3.25. Bytes 224-225, Vendor Specific Information.....	76
113	6.2.3. Non-Volatile Memory Specification	77
114	6.3. Cable Electrical Performance	77
115	6.3.1.1. Maximum Insertion Loss	78
116	6.3.1.1.1. 8.0 GT/s.....	78
117	6.3.1.2. Minimum Return Loss	79
118	6.3.1.2.1. 8.0 GT/s.....	79
119	6.3.1.3. Integrated Crosstalk Noise (ICN)	79
120	6.3.1.3.1. Multiple Disturber Near-End Crosstalk (MDNEXT).....	80
121	6.3.1.3.2. Multiple Disturber Far-End Crosstalk (MDFEXT).....	80
122	6.3.1.3.3. Weighting Functions.....	80
123	6.3.1.3.4. Integrated Crosstalk	81
124	6.3.1.3.5. ICN Limit.....	81
125	6.3.1.4. Maximum Pair-to-Pair Skew.....	82
126	6.3.1.5. Maximum Intra-Pair Skew	82
127	6.3.1.5.1. 8.0 GT/s.....	82
128	6.3.1.6. Maximum DC Resistance for Auxiliary Signals	82
129	6.3.1.7. Differential-to-Common Conversion Loss	82
130	6.3.1.8. CMI Line Capacitance	83
131	6.3.1.9. HiPot Requirement	83
132	6.3.1.10. Shielding.....	83
133	6.4. Cable Assembly Labeling	83
134	7. OTHER CONSIDERATIONS	84
135	7.1. Discovery of PCI Express Topology	84
136	7.2. Memory Allocation	84
137	7.3. Completion Timeouts of PCI Express Transactions	85
138	7.4. Bus Performance and Throughput for Interfaces with Long Flight Times	85
139	7.5. Interoperability with SAS cables	85
140	APPENDIX A. ACKNOWLEDGEMENTS	87
141		

142

List of Figures

143	Figure 1-1.	Connector for Fixed-Side and Free-Side	14
144	Figure 1-2.	Cable Interface	15
145	Figure 2-1.	CBLPRSNT# Signaling Diagram	18
146	Figure 2-2.	CADDR Signal Diagram	20
147	Figure 2-3.	CINT# Signal Implementation	20
148	Figure 2-4.	Example CMI Controller for Four x4 Cable Ports	21
149	Figure 2-5.	Device Address Map for a CMI Link	22
150	Figure 2-5.	Cable Management Interface Connectivity of Four Independent Cables	24
151	Figure 2-6.	Cable Management Interface Connectivity of Three Independent Cables, with Limited	
152		Cable Monitoring Capabilities	24
153	Figure 2-7.	Cable Management Interface Connectivity of three Independent Cables, with limited	
154		Cable Monitoring Capabilities	25
155	Figure 2-9.	Filter Implementation Example	29
156	Figure 3-1.	Electrical Parameter Allocation	32
157	Figure 3-2.	Transmitter Compliance Eye	36
158	Figure 3-3.	Receiver Compliance Eye	38
159	Figure 5-1.	SFF-8644 Fixed-Side Connector in Fixed-Side Cage	59
160	Figure 5-2.	Four Fixed-Side x4 Connector Configuration	60
161	Figure 5-3.	Temp Image, Connector Orientation Diagram	60
162	Figure 6-1.	Cable Assembly with SFF-8644 Connectors	63
163	Figure 6-2.	Structure of the PCI Express Cable CMI Memory Map	66
164	Figure 6-3.	Bytes 108-109, Propagation Delay	72
165	Figure 6-4.	Cable Assembly Insertion Loss	78
166	Figure 6-5.	Insertion Loss at F_{IL_ICN}	82
167			

List of Tables

169	Table 2-1.	DC Specifications for Auxiliary Signals	27
170	Table 2-2.	AC Specifications for Auxiliary Signals	28
171	Table 2-3.	PWR and MGTPWR Requirements	29
172	Table 3-1.	Jitter Budgeting Assumptions for Separate Clocks.....	32
173	Table 3-2.	Allowable Interconnect Lane-to-Lane Skew	33
174	Table 3-3.	Transmitter Path Compliance Eye Requirements at 2.5 GT/s.....	34
175	Table 3-4.	Transmitter Path Compliance Eye Requirements at 5.0 GT/s.....	35
176	Table 3-5.	Transmitter Path Compliance Eye Requirements at 8.0 GT/s.....	35
177	Table 3-6.	Receiver Path Compliance Eye Requirements at 2.5 GT/s.....	36
178	Table 3-7.	Receiver Path Compliance Eye Requirements at 5.0 GT/s.....	37
179	Table 3-8.	Receiver Path Compliance Eye Requirements at 8.0 GT/s.....	37
180	Table 3-9.	Free-Port Receiver Path Sensitivity Requirements at 8.0 GT/s	39
181	Table 3-10.	Fixed-port Receiver Path Sensitivity Requirements at 8.0 GT/s	39
182	Table 4-1.	Fixed-Side Memory Map	41
183	Table 4-2.	Sideband Messages.....	43
184	Table 4-3.	Fixed-Side Capability Registers	46
185	Table 4-4.	Cable Port Transmitter and Receiver Starting Values	53
186	Table 5-1.	Fixed-Side Connector Pinout	58
187	Table 5-2.	Fixed-Side Multi-connector Lane	59
188	Table 6-1.	Connector Pin assignments	63
189	Table 6-2.	Cable Wiring.....	63
190	Table 6-3.	Page 0 Lower Memory Fields	67
191	Table 6-4.	Page 0 Upper Memory Fields	68
192	Table 6-5.	Status Indicators	70
193	Table 6-6.	Lower Bound, Upper Bound, Base, and Span Values.....	72
194	Table 6-7.	Cabled PCI Express Capabilities 1	73
195	Table 6-8.	Cabled PCI Express Capabilities 2	73
196	Table 6-9.	Cable Technology Type	75
197	Table 6-10.	Non-volatile Memory Specifications.....	77
198	Table 6-11.	Cable Assembly Differential Characteristics Summary	77
199	Table 6-12.	Calculate ICN, MDNEXT and MDFEXT	79
200			
201			

1. Introduction

This is a companion specification to the *PCI Express Base Specification*. The primary focus of this specification is the implementation of cabled PCI Express[®]. No assumptions are made regarding the implementation of PCI Express-compliant Subsystems on either side of the cabled Link (PCI ExpressCard Electromechanical (CEM), ExpressCard™, ExpressModule™, PXI Express™, or any other form factor). Such form factors are covered in separate specifications.

1.1. Terms and Acronyms

x1, x2, x4, x8, x12, x16	x1 refers to one PCI Express Lane of Basic bandwidth; x4 to a collection of four PCI Express Lanes; etc.
Active Cable Assembly	A cable assembly that uses active components to buffer or condition the PCI Express signals.
Auxiliary signals	Signals not required by the PCI Express architecture, but necessary for certain desired functions or system implementation; for example, the CBLPRSNT# signal.
Box	An entity with an Upstream subsystem, one or more Downstream systems, or both, and contains a local power supply and other circuits that may include more than PCI Express devices and functions,
Cable Management Interface, CMI	The 2-wire management interface for communication with the cable and between Subsystems.
Cable Aggregation	The ability of a port to combine multiple cable assemblies inserted into the external connector to create a port larger than any of the individual connectors.
Cable Link Partner, CLP	The Fixed-Side Subsystem that a device and port are connected to via a cable
Connector	The portion of the mating interface that physically connects the Free-Side and Fixed-Side. The connector conforms to the mechanical definitions contained and contains the signal pinout defined herein. Unless otherwise stated, the definition includes four lanes.
Device	A component on either end of a PCI Express Link.

Downstream	<ol style="list-style-type: none"> 1. The relative position of an inter-connect/system element (Port/component) that is farther from the Root Complex. Within the context of this specification also referred to as <i>Downstream Subsystem</i>. 2. A direction of information flow where the information is flowing away from the Root Complex.
Downstream Facing Port, DSF	A cabled port in the Upstream Subsystem that links over the cable to a Downstream Subsystem.
Endpoint	A device with a Type 00h Configuration Space header.
Fixed-Side	The Subsystem a cable plugs into. Fixed-Side refers to either the Upstream or Downstream cable port; the term itself does not distinguish which end of the topology.
Free-Side	This is the cable side of a port. It is the removable side, or free side, of a cable connection.
Form Factor	In the context of this specification, form factor refers to other specifications that could be used as the foundation for implementing an external cabled PCI Express Port (e.g., PCI ExpressCard Electromechanical, ExpressCard, ExpressModule, Compact PCI Express).
G3	Global Power State in which no external power is provided to the Box. Refer to the <i>Advanced Configuration and Power Interface Specification</i> .
Host System	The Box which contains the PCI Express Root Complex for the PCI Express hierarchy. In the context of this specification, the Host System is an Upstream Subsystem with one or more DSF ports.
Hot-Plug	Insertion and/or removal of a cable and Downstream Subsystem into/from an active Upstream Subsystem.
Lane	One PCI Express Lane contains a differential pair for transmission and another differential pair for reception. A xN, Link is composed of N Lanes.
Link	A collection of one or more PCI Express Lanes providing the communication path between an Upstream and Downstream Port.
Legacy	Reference to the cable and parameters specified in the first two revisions of this specification that support 2.5 GT/s and 5.0 GT/s.
Local	A Fixed-Side port and inserted cable connector that identifies which end of the cable link is being referenced but is independent of upstream and downstream references. The CLP is at the opposite end of the cable assembly.
Multi-master Controller	A device that implements the Multi-master functions, with arbitration, specified in the <i>I2C Bus Specification</i> .
Passive Cable Assembly	A cable assembly that does not use any active components for the PCI Express signals. The cable assembly must still contain a memory device for system configuration information.
Port	<ol style="list-style-type: none"> 1. A group of Transmitters and Receivers located on the same device that define a Link when active. 2. Physical connector for a Fixed-Subsystem related to a link, typically referring to the Fixed-Side.
Port Bifurcation	The ability of an Upstream Subsystem to configure an external connector as multiple smaller ports.

Set	<u>A bit is Set when its value is 1b.</u>
Sideband Messages	A method for signaling events and conditions using the 2-wire interface between two linked components.
Subsystem	In the context of this specification, Subsystem is a generic term identifying either an Upstream or Downstream Fixed-Side containing the PCI Express device, CMI controller, and other circuits providing a cabled PCI Express Port.
Upstream	<ol style="list-style-type: none"> 1. The relative position of interconnect/system element (Port/component) that is nearest to the Root Complex. Within the context of this specification also referred to as <i>Upstream Subsystem</i>. 2. A direction of information flow where the information is flowing toward the Root Complex.
Upstream Facing Port, USF	A cabled port in the Downstream Subsystem that links over the cable to an Upstream Subsystem.
Vendor Specific Message	A message sent across the cable over the Cable Management Interface that is defined by the vendor. This is not related to in-band vendor-defined Messages allowed in the protocol.
Wake	A mechanism used by a component to request the reapplication of main power when in the L2 Link state. Two such mechanisms are defined in the <i>PCI Express Base Specification</i> : Beacon (using in-band signaling) and WAKE# (using side-band signaling).

1.2. Reference Documents

Following is a list of support documentation that should be referenced:

- ❑ *PCI Express Base Specification*, Revision 3.1
- ❑ *PCI Express OCuLink Specification*, Revision 1.0
- ❑ *PCI Hot-Plug Specification*, Revision 1.1
- ❑ *PCI Standard Hot-Plug Controller and Subsystem Specification*, Revision. 1.0
- ❑ *EIA 364 Series, Electrical Connector Test Procedures Including Environmental Classifications with Test Procedures*
- ❑ *EIA 364-1000, Environmental Test Methodology for Assessing the Performance of Connectors and Sockets used in Business Office Applications*
- ❑ *SFF-8644, Mini Multilane 12 Gbs 8/4x Shielded Connector*, Revision. 2.9
- ❑ *SFF-8636, Common Management Interface*, Revision. 2.6
- ❑ *I2C-Bus Specification* Revision 2.1
- ❑ *Advanced Configuration and Power Interface Specification*, Revision 6.1
- ❑ *SFF-8449, Shielded Cables Management Interface for SAS*, Revision 2.0.

1.3. Documentation Conventions

Following are a list of conventions that have been used throughout this specification:

□ Capitalization

Some terms are capitalized to distinguish their definition in the context of this document from their common English meaning. Words not capitalized have their common English meaning. When terms such as *memory write* or *memory read* appear completely in lower case, they include all transactions of that type.

Register names, and the names of fields and bits in registers and headers, are presented with the first letter capitalized and the remainder in lower case.

□ Numbers and Number Bases

Hexadecimal numbers are written with a lower case “h” suffix (like FFFh and 80h).

Hexadecimal numbers larger than four digits are represented with a space dividing each group of four digits, as in 1E FFFF FFFFh. Binary numbers are written with a lower case “b” suffix (1001b and 10b). Binary numbers larger than four digits are written with a space dividing each group of four digits, as in 1000 0101 0010b.

All other numbers are decimal.

□ Implementation Notes

Implementation Notes should not be considered a part of this specification. They are included for clarification and illustration only.

1.4. Specification Contents

This specification contains the following information:

- Subsystem requirements
- Sideband signaling and usage models
- Cable hot insertion and removal
- Subsystem electrical budgets
- Cable and connectors electrical budget
- Cable and connectors specifications
- Power provisioning

This specification describes the new Third Generation of the external cabling interface specifications and covers the new connectors, cables, signaling, Cable Management Interface, and sidebands.

1.5. Objectives

The objectives of this specification are as follows:

- ❑ Define PCI Express external cables and associated connectors
- ❑ Support PCI Express data rates of up to 8.0 GT/s, while maintaining compatibility with the 5.0 GT/s and 2.5 GT/s signaling rate specifications
- ❑ Support standard PCI Express components, as defined by the *PCI Express Base Specification*
- ❑ Forward looking for future scalability
- ❑ Maximize cable interoperability for user flexibility
- ❑ Enable Hot-plug as a native function
- ❑ Allow revolutionary partitioning of the PC architecture
- ❑ Upgradeability
- ❑ Enable the construction of adapters to provide compatibility with legacy downstream subsystems

1.6. Overview

PCI Express is the third generation of a multi-purpose I/O interface that is used across the computing industry (from mobile through high-end servers and communication equipment.) The broad usage and versatility of this technology allows for system extensions to external input/output Subsystems that meet the needs for specific target applications and/or environments.

Cabled PCI Express targets applications, including but not limited to:

- ❑ Split-systems, or disaggregate PCs, with a desktop “console” that contains removable media drives (e.g., CD/DVD), memory modules, I/O ports (e.g., USB, IEEE-1394), and audio jacks
- ❑ I/O expansion to extend the I/O card capabilities of the main system for support of different form factors, including legacy Subsystems, test and measurement, and instrumentation equipment
- ❑ Server expansion I/O to support conventional PCI Express add-in cards (with or without Hot-Plug support) and/or ExpressModules
- ❑ Location of the graphic Subsystem (i.e., controller and memory) external to the main systems unit

System-level support for cabled PCI Express is possible through implementation on expansion card (on a PCI Express Card Electromechanical, ExpressModule, PXI Express, ExpressCard, or direct from a system board).

This specification defines connectors and implementation for supporting up to x16 cabled PCI Express Links.

Utilizing off-the-shelf PCI Express components is the focus of this specification, while providing enough flexibility for design and implementation of dedicated components for driving the cabled interconnect. As such, much of this specification builds upon other PCI Express form factor specifications. However, not all components inherently function correctly in this application.

New to this version of the specification are electrical and mechanical requirements. These requirements may introduce some backward compatibility concerns.



Note: This revision of the specification calls for a new connector family, one commonly known as *SFF-8644*, that is not mechanically compatible with the connectors of the previous versions of this specification. Therefore, cable ports and assemblies designed to this specification are not directly compatible with ports and assemblies designed specifically to meet the previous revisions of this specification. The new connector scheme provides for aggregation of connectors and cable assemblies to support Links widths of x16, as well as the option for Port Bifurcation, to support multiple links to a single Upstream physical port/connector. See Figure 1-1.




Figure 1-11-1. Connector for Fixed-Side and Free-Side

New features enabled by this revision of the specification include:

- ❑ Bifurcation of cable ports, which enables Downstream facing ports with more than one x4 connector to be split into smaller, separate links. (Optional)
- ❑ Cable Aggregation, which allows multiple smaller cable widths to create a larger port. (Optional)
- ❑ Sideband Messages, a new mechanism to send sideband signals over a serial protocol.
- ❑ Vendor Specific ~~Communication~~communication, a mechanism for sending Vendor Specific Messages outside of the PCI Express signal path. (Optional)
- ❑ Cable Monitoring, which allows for monitoring of active cable component status. (Optional)
- ❑ Fiber Optic Support (Optional)
- ❑ Compatibility with SFF-8449 cables. These cables are to be used to create x4 links and do not have access to several of the above features.

Previous versions of this specification utilized/employed a fixed-size component for each supported link width (e.g., a x1 link had specific x1 cable ports and a x1 cable assembly, a x4 link had a specific x4 cable port and x4 specific cable assembly, up through x16 link widths). Unlike some other form-factors, components of various link widths were not interoperable. Introduced in this revision of this specification, the SFF-8644 connectors are defined for increased performance characteristics and new features. The connector family scales a link from x4 connection up to x16. A unique x1 connector is not provided and instead the new x4 connector is used for any size link up to, and including, x4 link. The x8 link aggregates two connectors, and x16 aggregates four connectors. This allows for a more universal and scalable cable port.

The introduction of 8.0 GT/s in the *PCI Express Base Specification*, requires some foreknowledge of the physical link interconnect electrical performance for the purpose of configuring equalizer starting points. Since the performance of a cable is unknown to either Fixed-Side Subsystem, the use of the CMI and cables with memory has been employed for storing and retrieving cable assembly performance information (Figure 1-2 shows the cable interface).

The sideband messaging scheme changes such that individual wires are no longer used for each sideband signal. This revision of the specification also enables an optical friendly architecture for those needing either distance or electrical isolation between Subsystems.




Figure 1-21-2. Cable Interface

330

2. Auxiliary Signals

Auxiliary signals are provided on the connector to assist with system-level functionality or implementation. The high-speed differential signaling levels are compatible with advanced silicon processes while all low-speed Auxiliary signals are defined to be compatible with +3.3 V signaling unless otherwise noted.

335

Besides the signals that are required to transmit/receive data on the PCI Express interface, there are also signals that may be necessary to implement the PCI Express interface in a distributed system environment, or to provide certain desired functions, including power for data conditioning within the connector backshell. These signals are referred to as the auxiliary signals.

340

The PCI Express cable connector and cabling support the following Auxiliary signals:

- ❑ **PWR:** Power provisioning to the connector backshell is provided to allow for signal conditioning components within the cable assembly. A wire must not be provided within the cable.
- ❑ **MGTPWR:** Power supplied to the connector backshell for cable management components that are needed while the link is not active. This needs to be active if the Subsystem has power. A wire must not be provided within the cable.
- ❑ **CBLPRSNT#:** Cable present detect, an active-low signal pulled-down by the Free-Side when it is inserted into the Fixed-Side Connector. A wire must not be provided within the cable.
- ❑ **CADDR:** This signal is used to configure the Upstream cable management device address. A wire must not be provided within the cable.
- ❑ **CINT#:** This signal is asserted by the cable assembly to indicate a need for service via the Cable Management Interface controller. A wire must not be provided within the cable.
- ❑ **CMISDA:** Management interface data line. Used for both initial link setup and sideband messages when used with proper cables.
- ❑ **CMISCL:** Management interface clock line. Used for both initial link setup and sideband messages when used with proper cables.

345

350

355

2.1. PWR

Power for the optional active cable circuitry within the cable backshell is required. This pin may not be powered in all Subsystem power states such as when the Fixed-Side Subsystem is in a power saving state.

Due to the low level signaling of the PCI Express interface, it is strongly recommended that sufficient decoupling of all power supplies be provided. This is recommended to ensure that power supply noise does not interfere with the recovery of data from a remote Subsystem. Some basic guidelines to help ensure a quiet power supply are provided in Section 2.9.



Note: The details given in this section are guidelines only. It is the responsibility of the designer to properly test the design to ensure that system board circuitry does not create excessive noise on power supply or ground signals at the connector interface.

To protect system operation, the PCI Express external cable assemblies (that operate from board-supplied power) must follow the requirements listed in Table 2-3 during Hot-Plugging and normal operation. The requirements for the power supply voltages are defined at the Fixed-Side connector.

2.2. MGTPWR

MGTPWR to the cable assembly is required to enable cable management circuitry. It is recommended that MGTPWR remain within operational limits when in power states other than G3 to enable remote power control and decrease cabled link enable times. Connectors allow MGTPWR to energize simultaneously with the sideband signals. Devices in the cable must be tolerant of this condition.



Note: MGTPWR is removed and re-applied to reset the cable logic.

SFF-8644/8614 has the same length pins for MGTPWR, PWR, CADDR, CINT#, CMISDA, and CMISCL. This means that there is no guarantee that the memory device on the cable will be powered before other signals are applied. You must specify these devices to be tolerant of this condition (i.e., so the devices do not suffer latch-up conditions).

2.3. CBLPRSNT#

A cable assembly presence detection mechanism is specified, through CBLPRSNT#, to indicate the attachment of a cable assembly to a Fixed-Side port. The Free-Side must assert this signal low to indicate cable attachment. It is recommended that this signal not be directly grounded, but rather asserted with a pull-down or active driver in which the value should be determined by cable manufacturer based on RP_RUX value specified and the voltage requirements specified in Table 2-1. The assertion of CBLPRSNT# provides no information about the far end of the cable, including whether or not it is inserted into the Fixed-Side of a CLP.

Upon detection of CBLPRSNT# asserted, the Fixed-Side management controller must read the cable memory's first 256 bytes to obtain cable information and begin monitoring of CINT#. See to Section 4.2.1 for specific usage in power sequencing. The Fixed-Side must provide a pull-up resistor, of value R_{P_AUX} , on CBLPRSNT# to MGTPWR to passively de-assert CBLPRSNT#.




Figure 2-12-1. CBLPRSNT# Signaling Diagram



Observation: CBLPRSNT# being asserted by the cable determines if the cable is present. This can be used as a link disable by the cable assembly by de-asserting CBLPRSNT# in an active fashion. This may be useful for some cable assemblies, but specific implementations are beyond the scope of this specification.

2.4. CADDR

The CADDR signal enables the use of a programmable address pin for the Free-Side CMI. Cable assemblies using SFF-8644 connectors at both ends use this pin as the address reference for the Least Significant Bit, or LSB, of the Free-Side memory address. The Free-Side must tie the LSB of the 7-bit address low with a resistor that allows the signal to remain below V_{OL_SB} when the Fixed-Side is open and above V_{OH_SB} when the Fixed-Side has a pull-up of R_{P_AUX} . The address can be changed by the Fixed-Side by overdriving the weak pull-down of the Free-Side. For cable assemblies supporting side-band messages, address A0h must contain the cable description information, while duplicate information at address A2h is optional. If there is a single memory device, CADDR may be ignored within the cable assembly such that the address is fixed to address A0h. However, cable assembly status information, such as the interrupts and monitors described in Section 6.2.1, are only be available to the USF Fixed-Side. Upstream Fixed-Side management controllers must read from address A0h for configuration data if *acknowledge* is not received when attempting to read from address A2h. If the Free-Side can respond to requests to address A2h, then the cable description information must also reside in the device at A2h.



Note: Addresses for the Cable Management Interface follow industry convention of specifying address in 8-bit notation, while the actual addresses are 7-bit and an additional bit, in the LSB position, indicates a Read or Write operation. For example, a write to the Free-Side's memory location is A0h and a read to the same location is A1h.



Observation: Upstream Fixed-Side management engine is permitted to listen for configuration information at address A0h. This may decrease configuration time.

If a cable assembly supports a single memory device, configuration information is only available if the relevant end of the cable is powered.

The Upstream Fixed-Side must drive CADDR high, via resistor of value R_{P_AUX} , or active driver.

Support of a cross-link configuration requires an active driver on the Upstream Fixed-Side. Once the PCI Express protocol has determined which portion of the link is the Upstream Subsystem, the Fixed-Side Subsystem that has been determined to be the Downstream Subsystem must allow the CADDR to float or drive the signal low to configure the Downstream end of the cable for management purposes. In the cross-link system, beginning with initialization, there are times when both cable memory locations are at the same address. Upstream Subsystems that do not support cross-link are permitted to implement a passive solution.

The Downstream Fixed-Side must not connect to CADDR. The Downstream end of the cable uses the default address of address A0h for the cable memory with CADDR not connected by the Fixed-Side. Figure 2-2 shows an example of a common memory device configuration within a cable assembly. In the management device within the Free-Side, pin A0 is tied low with a resistor that allows CADDR to drive the signal high by the upstream Fixed-Side. Pins A1 and A2 are tied low.

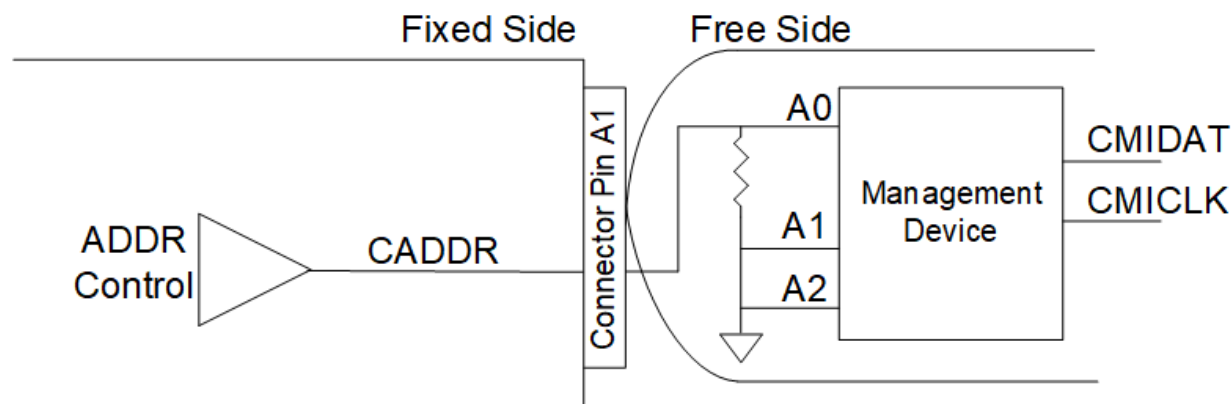


Figure 2-22-2. CADDR Signal Diagram

2.5. CINT#

430 CINT# enables the use of a Free-Side driven interrupt to indicate a need for service via the Cable Management Interface. CINT# is localized to the near end Free-Side of the cable only and does not provide any indication across the cable to the opposite Fixed-Side. The Upstream Fixed-Side receives only the CINT# from the end of the cable inserted to the Upstream Free-Side and then reads address A2h, and the Downstream Fixed-Side receives only the CINT# from the Downstream Free-Side and reads address A0h. A CMI Controller receives no indication of the interrupt being asserted by the cable assembly's opposite end Free-Side to the Fixed-Side of the CLP. Transferring any indicators from one end to the other is beyond the scope of this specification.

435

The Free-Side requests interrupt service by continuously asserting the CINT# pin until the Fixed-Side management controller clears the interrupt flag of the Free-Side memory.

440 The Fixed-Side must provide a pull-up resistor, of value R_{P_AUX} , on CINT# to MGT_PWR to passively de-assert CINT#.




Figure 2-32-3. CINT# Signal Implementation

2.6. Cable Management Interface

The Cable Management Interface, or CMI, is a two-wire interface through which various cabled system components communicate with each other. This type of interface is common for Fixed-Side Subsystems to read information from cable assemblies. It is a private interface meant to only communicate between the cabled Subsystem's management interface components. Multiple devices within a Fixed-Side Subsystem must not share CMISCL and CMISDA. A separate interface must be implemented to interface with the PCI Express devices and any other system member requiring information from the CMI. Figure 2-4 illustrates an example of a CMI controller implementation of four physical cable ports.

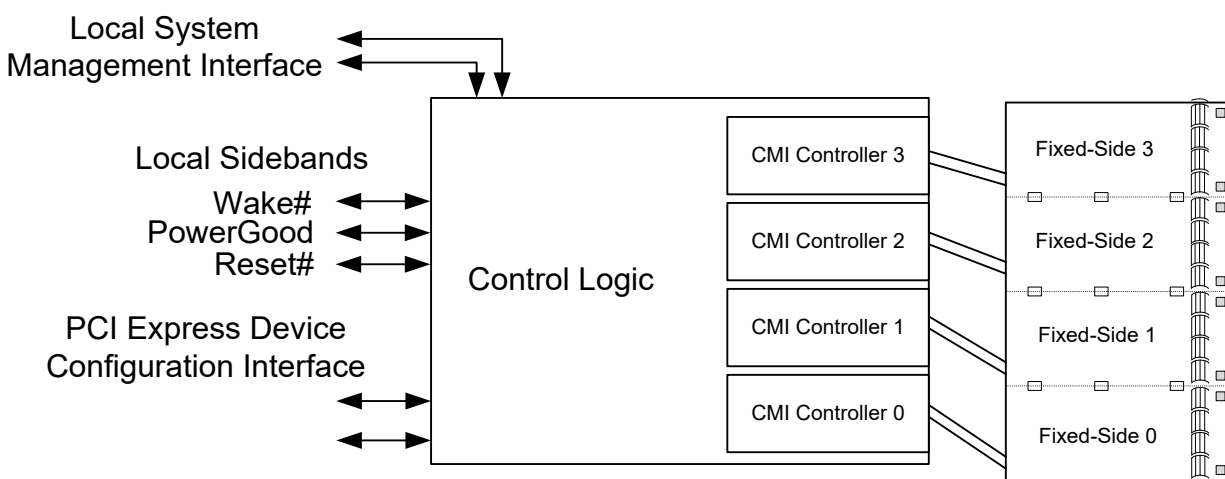


Figure 2-4. Example CMI Controller for Four x4 Cable Ports



Implementation Note: The CMI may be implemented in programmable logic, micro-controller, or CPU. The specific implementation may be dictated by the features to be supported by the Fixed-Side Subsystem.

Using the Cable Management Interface, each Fixed-Side of the link reads information from the cable assembly to configure the PCI Express interface. Optionally, CMI enables the sending and receiving of sideband messages between the two Fixed-Side Subsystems, located at addresses A4h and A6h, if the cable assembly is capable, as indicated by the support for Sideband Messages bit in the Cabled PCI Express Capabilities 1 register of the Free-Side. [A map of the devices addresses for a CMI link is shown in Figure 2-5.](#)




Figure 2-5. Device Address Map for a CMI Link

Each Fixed-Side Subsystem must implement a multi-master controller to enable communication between the Fixed-Side Subsystems. [Figure 2-1 shows the addresses for each device present on a CMI link.](#) A list of messages that represent sideband signals is provided in Table 4-2. In addition, Vendor-Specific Messages may also be sent via the Cable Management Interface if both the cable and devices support this operation. To support long cables the bus must operate at 100 kHz as the default data rate.

The CMI must support a clock stretching timeout specified by $T_{\text{CMI_CLK_HOLD}}$ in Table 2-2.

With the Cable Management Interface, a device in the cable assembly can provide manufacturer information, reports certain errors, accepts control parameters, and returns current status.

The Cable Management Interface conforms to the Management Interface Protocol described by the *PCI Express OCuLink Specification*. The memory maps are described in Section 6.2.1.

The Fixed-Side Subsystems provide pull-ups, of value $R_{\text{PU_CMI}}$, for CMISCL and CMISDA to the MGTPWR rail, and the components of these signals need to have a 3.3 V signal tolerance. Cable loading limits are specified in Table 2-2. Additional specific parameters for the devices on the CMI interfaces are listed in Table 2-1 and Table 2-2. The falling edge of the waveform must remain below $V_{\text{IL_CMI}}$ after the incident edge until the signal is to transition to a value of one.



Observation: Process technologies have allowed significant decreases in the transition time of driver outputs, even for 2-wire interfaces such as CMI. To allow fast edges to properly transition on the falling edge, impedance matching of the CMI device buffer to the $Z_{\text{CMI_PCB_ZO}}$ impedance target of the Fixed-Side may be required. Some devices may be excluded from this application, including devices with fast edges but high impedance.

The memory map for the storage device in the cable assembly is described in Section 6.2.2 which provides the information stored in the cable assembly that is available to the system. The information needed by the Upstream and Downstream Subsystems for configuration includes:

- ❑ Cable attenuation to calculate the starting points for both transmit and receive equalizers
- ❑ [PCI Express specific cable and bit rates supported by the cable](#) ~~Bit rates supported by the cable~~
- ❑ Number of Lanes supported by the cable
- ❑ Flight time of the cable
- ❑ Cable technology

This information is used to configure the Fixed-Side devices, both logically and electrically. It should be noted that further fine tuning of the transmitters and receivers occurs during the link equalization procedure that is specified in the *PCI Express Base Specification*.

Power Sequencing in Section 4.2.1 describes when the Free-Side memory device is read. The CMI Controllers for each cable port must access the cable assembly memory devices, located at address A0h for Downstream devices and address A2h or address A0h for Upstream devices. Physical cable port 0 is the only port required to be accessed if the Free-Side supports more than a single x4 Link. The Cabled PCI Express Capabilities 1, Cabled PCI Express Capabilities 2, Cable Technology, and Attenuation at Frequency registers are used to configure the PCI Express device logical parameters, equalization configuration, and are used for some system parameters. The communication process to configure the PCI Express device and other system parameters, once the information is retrieved and determined, is implementation-specific.

For a sideband enabled cable assembly, the Cable Management Interface transmits status change events at the system level. See Section 4.1 for a list of status registers that are read and messages that are received.



Implementation Note: Not all cable assemblies have bidirectional communication for system status communication via the cable management interface. The implication is that the external device may not be ready when the Host System begins accessing the PCI Express bus. It is recommended that a visible indicator ~~is-be~~ implemented to alert the user that the external device is ready for the system to begin configuration. However, details of such an indicator are beyond the scope of this specification.

To minimize the impact of topology, loading, and addressing, there is a limit of four loads on the Cable Management Interface. There is a controller for each Fixed-Side and two devices in the cable. The Fixed-Side must provide a CMI controller for each logical link that may be created, or for each physical cable port that is to be monitored for cable status.

For Upstream or Downstream Fixed-Sides that support either Cable Aggregation or Port Bifurcation, a controller for each x4 interface to be concurrently supported per physical x4 connector is required for setting the respective equalizers. If a cable assembly is wide enough for the entire physical link, referred to as a monolith cable assembly, then only the Cable Management Interface associated with the lowest lane may be used for all the management and configuration information. For example, if the port supports sixteen lanes, does not support Port Bifurcation, does not support aggregation, then only the management controller associated with lanes 0-3 is required. The configuration information from the cable assembly is used to configure all lanes. Figure 2-5 shows the cable management interface connectivity of four independent cables.



Implementation Note: If a Fixed-Side port only implements a management controller for the connector associated with lanes 0-3, configuration data and sideband messages are supported for the link. However, active cable assembly status information for lanes 4 and above would not be available. If a multi-connector cable assembly only implements a Cable Management Interface for the connector associated with lanes 0 to 3, then active cable assembly status information for lanes 4 and above would not be available to the management controller on the Fixed-Side.

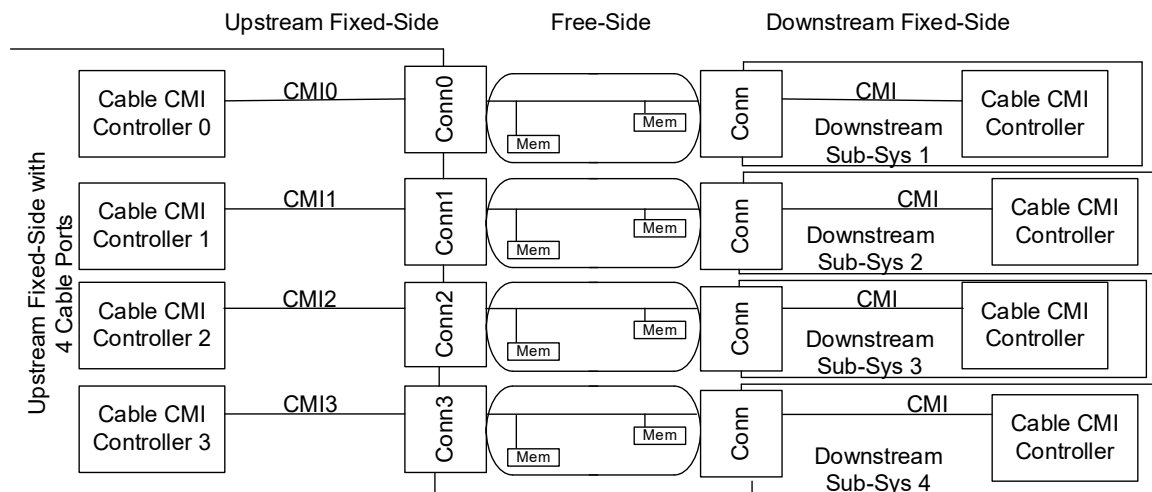


Figure 2-62-5. Cable Management Interface Connectivity of Four Independent Cables

Figure 2-6 illustrates a bifurcated port with two x4 links and one x8 link. The x8 link aggregates two x4 cable assemblies. In this case, there are four Upstream CMI controllers and four Downstream CMI controllers. This x8 Downstream Fixed-Side needs the CMI controller to support Cable Aggregation and read each cable memory device to access configuration information. However, sideband messages are only required on the CMI link associated with the lowest lane in the link (CMI for Downstream Subsystem 3). The CMI associated with the higher lanes may also send the sideband messages, but each Subsystem must reference the CMI associated with the lowest lane for interpretation of the sidebands for the linked Subsystem.

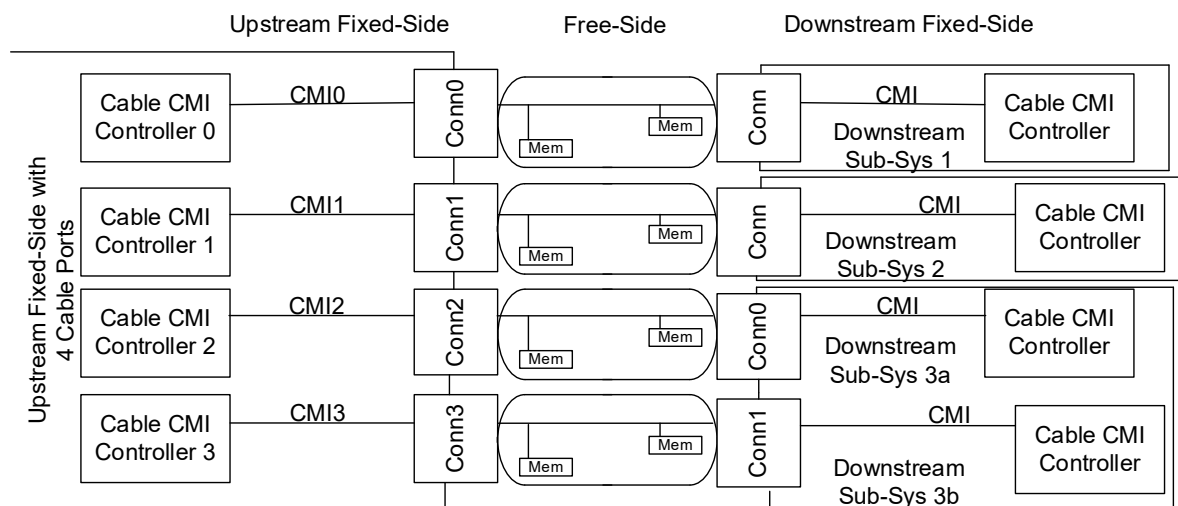


Figure 2-72-6. Cable Management Interface Connectivity of Three Independent Cables, with Limited Cable Monitoring Capabilities

Figure 2-7 is similar but shows a Downstream x8 Fixed-Side device that does not support Cable Aggregation and requires the use of x8 cable assembly. The x8 assembly only has a single cable management interface. The Upstream CMI controller must interpret sideband messages received on CMI2 to reference the entire x8 link (the upper two cable connectors). Cable assembly status information is not read from the assembly since neither the cable assembly, nor the Downstream Subsystem, have the components. If the cable assembly contained memory devices, the Upstream port would be able to read cable assembly status information while the Downstream Subsystem would not. The Upstream Subsystem would be able to read the Downstream cable assembly device if the upper quad's management bus traversed the cable assembly.

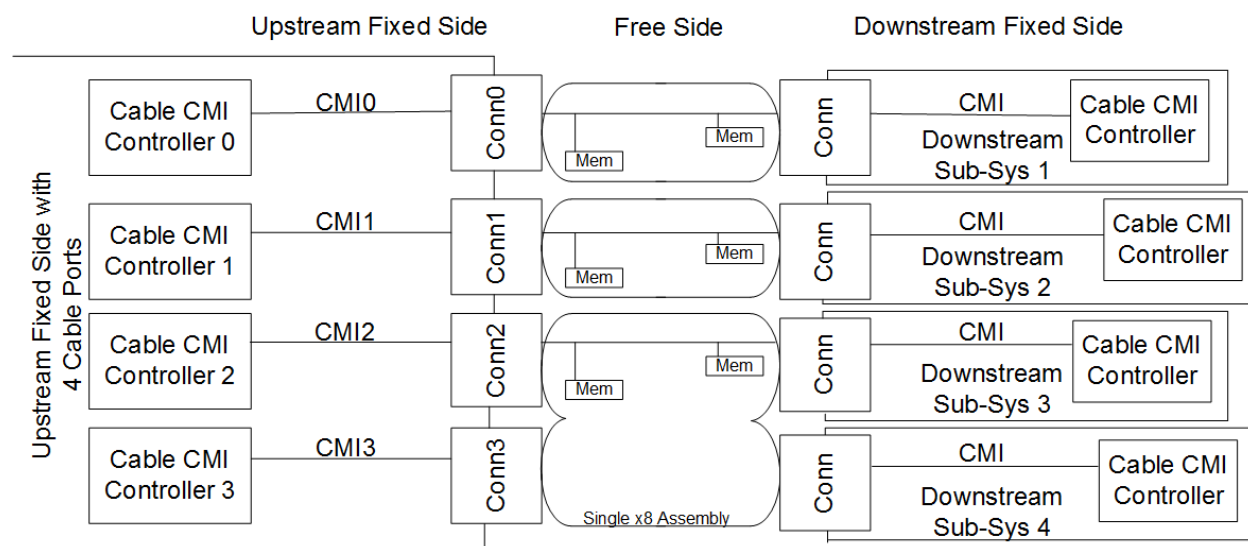


Figure 2-82-7. Cable Management Interface Connectivity of three Independent Cables, with limited Cable Monitoring Capabilities

2.6.1. Capacitive Load of High-power Cable Management Interface Lines

The capacitive load for a Fixed-Side is defined by C_{CMI_FIXED} and includes both the capacitance of the device and capacitance of the routing channel. The capacitive load for a cable assembly is defined by $C_{CMI_CABLE_LOAD}$.

2.6.2. Minimum Current Sinking Requirements for Cable Management Interface Devices

Devices are required to sink a minimum current of I_{SINK} while maintaining the V_{OL_CMI} (max) of 0.4 V.

2.6.3. Cable Management Interface *Back Powering* Considerations

565 Unpowered devices connected to either a low-power or high-power CMI segment must provide, either within the device or through the interface circuitry, protection against “back powering” the Cable Management Interface. Unpowered devices must meet I_{LEAK} parameters in Table 2-1.

2.7. Electrostatic Discharge

Electrostatic Discharge (ESD) requirements for auxiliary signals are identical to those specified in the PCI Express signal protection. See section 3.5 for ESD requirement details.

2.8. Auxiliary Signal Parametric Specifications

Figure 2-8 illustrates the locations where measurements are to be taken. Some measurements may require a test fixture at the test point for access to the signal and care should be taken to limit or remove the impact of the test fixture.

2.8.1. DC Specifications for Auxiliary Signals

Table 2-1 lists the DC specifications for Auxiliary signals

Table 2-12-1. DC Specifications for Auxiliary Signals

Symbol	Parameter	Min	Max	Units	Notes
RP_AUX	Resistor for	9 k	10 k	Ω	2
RPU_CMI	Pull-up Resistor for Fixed-Side CMI signals	2.13 k	2.17 k	Ω	3,5,6
VOH_SB	Voltage output high for cable auxiliary signals	$V_{MGT PWR}-0.5$	$V_{MGT PWR}+0.3$	V	1,2
VOL_SB	Voltage output low for cable auxiliary signals	0	0.4	V	1,2
VIH_SB	Voltage input high for cable auxiliary signals	$V_{MGT PWR} * 0.7$	$V_{MGT PWR}+0.5$	V	1,2
VIL_SB	Voltage input low for cable auxiliary signals	-0.3	$V_{MGT PWR} * 0.3$	V	1,2
VOH_CMI	Voltage output high for CMI signals	$V_{MGT PWR}-0.5$	$V_{MGT PWR}+0.3$	V	1,3
VOL_CMI	Voltage output low for CMI signals	0	0.4	V	1,3,6
VIH_CMI	Voltage input high for CMI signals	$V_{MGT PWR} * 0.7$	$V_{MGT PWR}+0.5$	V	1,3
VIL_CMI	Voltage input low for CMI signals	-0.3	$V_{MGT PWR} * 0.3$	V	1,3,6
ISINK	Current a CMI device must be able to sink	3.0		mA	
I _{LEAK}	Input Leakage Current of Active Cable Components	-5	+5	μ A	
CCMI_CABLE_LOAD	Total capacitance of each CMI signal in a Cable Assembly	<330		pF	7

Notes:

1. Measured at the Fixed-Side Subsystem connector.
2. Applies to CBLPRSNT#, CINT# ~~and CADDR~~ Fixed-Side Pull-ups, and CADDR# Free-Side Pull-down.
3. Applies to CMISDA and CMISCL.
4. Capacitance includes devices and routing path.
5. Pull-up resistors are implemented by both Fixed-Sides, so two pull-up resistors in parallel are used to generate the high levels and the drivers must be able to meet low voltages with the lower equivalent resistance.
6. The value needed to balance cable capacitive limits and timings requires a 2.15 k 1% resistor tolerance.
7. Total capacitive load of each cable assembly's CMI signals, including components.

2.8.2. AC Specifications for Auxiliary Signals

Refer to Appendix B in the *PCI Express OCuLink Specification* for most of the CMI timing requirements. Table 2-2 lists additional requirements to those listed in the *PCI Express OCuLink Specification*.

Table 2-2-2. AC Specifications for Auxiliary Signals

Symbol	Parameter	Value	Units	Notes
F _{CMI_MAX}	CMI Maximum Operational Frequency	100	kHz	1
T _{CMI_CLK_HOLD}	Serial Interface Clock Hold-off (Clock Stretching)	<500	μs	2,3
T _{MGT_INIT}	Management Interface Initialization Time	<500	ms	4
T _{SB_TIMEOUT}	Time limit between sending CMI message update	<500	ms	
Z _{CMI_PCB_ZO}	Impedance of CMI PCB traces	55	Ω	5
Z _{CMI_CABLE_ZO}	Impedance of CMI Wires in Cable Assembly	>50	Ω	6
T _{CMI_SUPPRESS}	Width of noise spikes which an input filter must suppress on the CMI inputs	50	ns	
V _{HYSTERESIS}	Hysteresis of Schmitt trigger inputs	0.05* V _{MGTPWR}	V	
C _{CMI_FIXED}	Capacitance of CMI Interface for the Fixed-Side	30	pF	5,7

Notes:

1. See Section 2.6 for information on the behavior of F_{CMI_MAX}. This excludes clock stretching.
2. Amount of time the cable or CLP is permitted to hold CMISCL low before continuing with a read or write operation.
3. It is recommended that clock stretching by the controller be disabled if the Free-Side is optical.
4. SFF-8449 cables specify a T_{MGT_INIT} of 2000 ms.
5. Applies to CMISCL and CMISDA traces on Fixed-Side PCB.
6. Applies to CMISCL and CMISDA wires in the cable assembly.
7. Pull-up resistors implemented by both Fixed-Sides, so two pull-up resistors in parallel are used to generate the high levels and the drivers must be able to meet low voltages with the lower equivalent resistance.

580

2.9. Power Requirements

Power supply filtering per cable is recommended to minimize system noise from impacting the circuits in the cable assembly, or for one cable assembly impacting another. Table 2-3 lists the PWR and MGTPWR requirements.

Table 2-32-3. PWR and MGTPWR Requirements

Symbol	Parameter	Value	Units	Notes
V_{MGTPWR}	Management Power Pin Voltage	3.0-3.6	V	4
	Maximum Sustained Current	30	mA	
V_{PWR}	Power Pin Voltage	$3.3 \pm 5\%$	V	1, 4
	Power supply noise tolerance (Peak)	<50	mV	4
	Power Consumption for Free-Side	<1.5	W	3, 5
	Maximum Sustained Current 30 mA 1,2	<500	mA	1, 2
	Capacitive Load of the Cable Assembly <15 μF 1	<15	μF	

Notes:

1. This is per pin.
2. Current requirements apply to the current through the contacts.
3. Maximum power consumption of the cable assembly must not exceed 1.5 W from 500 ms after power up.
4. Includes noise and ripple
5. This is the total power limit for the module. How power is consumed per pin is beyond the scope of this specification.

585

An example of a filter implementation for the power pins per connector plug receptacle power pin is shown in Figure 2-9. It is provided for informational purposes only and does not imply a specific implementation. For this example, the filter inductor is 1 μH with an ESR of 100 m Ω . The bulk capacitors are 22 μF with an ESR of 220 m Ω .

590

It is strongly recommended that each power pin be filtered independently.

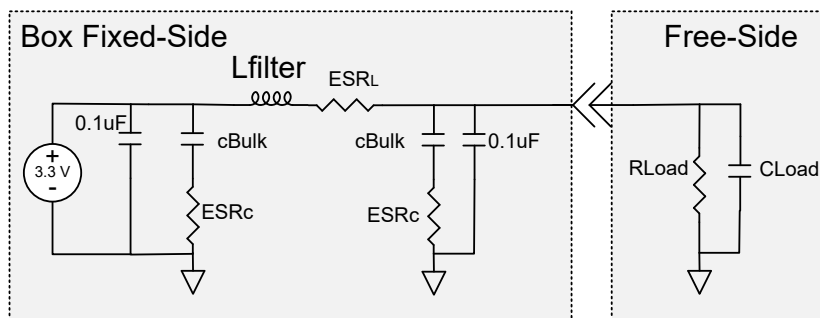


Figure 2-92-9. Filter Implementation Example

3. PCI Express Signals

This specification defines the requirements for PCI Express cables that support links using a PCI Express PHY. This specification does not support links using M-PCI Express (eg. M-PHY).

Each PCI Express Lane consists of a pair of differential signals. The Transmitter pair is labeled PETpN and PETnN where **N** is the lane number (starting with 0); **p** is the true signal and **n** is the complement signal. The Transmitter pair originates in the Fixed-Side Subsystem and is connected to a Receiver pair at the opposite Fixed-Side Subsystem at the opposite ~~side-end~~ of the cable. The Receiver pair is labeled PERpN and PERnN with the same labeling convention as the Transmitter pair.

Support for Polarity Inversion is required on all PCI Express Receivers across all Lanes independently. The **p** and **n** connections may be reversed to simplify trace routing and minimize vias if needed. All PCI Express Receivers incorporate automatic Polarity Inversion as part of the Link Initialization and Training and corrects the polarity independently on each differential pair. Refer to Chapter 4 of the *PCI Express Base Specification* for more information.

Support for Lane Reversal is an option on all PCI Express receivers. An example of Lane Reversal Lane 0 of an Upstream Port connected to Lane N-1 of a Downstream Port where either the Downstream or Upstream device may reverse the Lane order to configure a **xN** Link. Refer to Chapter 4 of the *PCI Express Base Specification* for more information.

Polarity Inversion and/or Lane Reversal must not be implemented within the cable assembly.

A *null* modem function for connecting the transmit pair on one ~~side-end~~ to the receive pair to the other ~~side-end~~ is performed within the cable assembly. External cable connector pin assignments are identical at both sides of the Link. The wiring of the cable assembly is shown in Table 6-2.

3.1. Interconnect

In the context of this specification, *interconnect* comprises everything between the pins of the Transmitter package and the pins of the Receiver package. This consists of traces on printed circuit boards, cable, AC-coupling capacitors, and connectors.

The interconnect total capacitance to ground seen by the Receiver Detection circuit (see Chapter 4 of the *PCI Express Base Specification*) must not exceed 3 nF, including capacitance added by attached test instrumentation. Note that this capacitance is separate and distinct from the AC-coupling capacitance value (see Section 3.2.1). The target impedance for the interconnect is 100 Ω .

3.1.1. Link Definition

Typical cabled PCI Express Links, from source to destination, consist of the following:

- Transmitters on an ASIC on a printed circuit board
- Package fan-in-out trace topologies
- PCB coupled microstrip and/or stripline traces
- AC-coupling capacitors
- Vias for layer changes
- Cable mated connector
- External raw cable
- Coupled microstrip line and/or stripline traces
- Receivers on an ASIC on a printed circuit board

Figure 3-1 shows the electrical parameters for the Link are subdivided into:

- Upstream Subsystem
- Cable including mated connectors
- Downstream Subsystem

Note that for validation purposes, the separation is somewhat different to facilitate ease of connecting any test and measurement equipment.

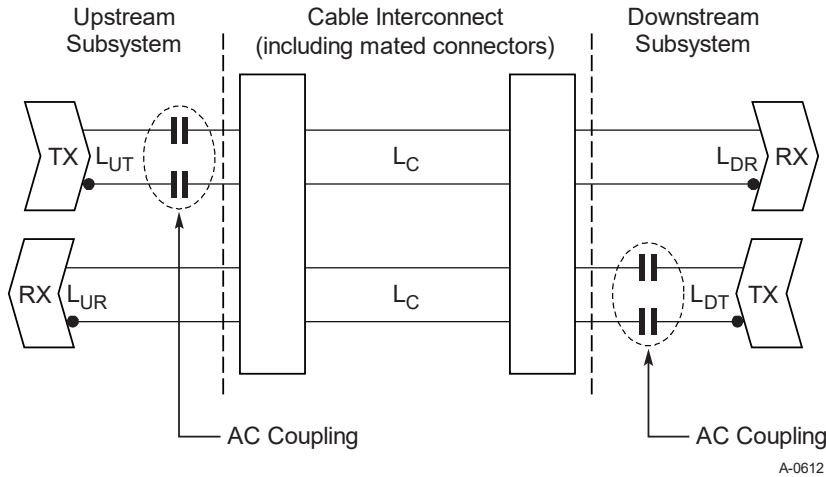


Figure 3-13-1. Electrical Parameter Allocation

3.2. Electrical Budget

3.2.1. AC-coupling

AC-coupling is required at the Transmitter Fixed-Side (see Figure 3-1) with values as specified in Chapter 4 of the *PCI Express Base Specification*. A passive cable assembly does not implement any AC-coupling capacitors. Active cables must provide AC-coupling capacitors in an implementation specific manor.



Note: SATA and SAS systems implement the capacitors on the Rx Subsystem and not the Tx Subsystem. With both specifications utilizing the SFF-8644 connector, it is possible to cross-plug a PCI Express port with a SAS port. It is strongly recommended that PCI Express receivers tolerate a DC connection to SATA or SAS transmitters without damage. The method used to achieve this is beyond the scope of this specification

3.2.2. Jitter from Clock Sources

A portion of the link budget is required for the differences in the clocks of the two fixed Subsystems. The jitter requirements for the Upstream and Downstream clocks that are referenced by the cable link are in Table 3-1.

Table 3-13-1. Jitter Budgeting Assumptions for Separate Clocks

Jitter Parameters	Symbol	Values (RMS in ps)	Notes
Link Jitter for 2.5 GT/s	T _{INDCLK LINK JITTER 2.5G}	4.30	1 2
Link Jitter for 5.0 GT/s	T _{INDCLK LINK JITTER 5G}	2.0	1, 2
Link Jitter for 8.0 GT/s	T _{INDCLK LINK JITTER 8G}	0.50	1, 2

Notes:

1. Defined for either SSC or non-SSC implementations.
2. Filter functions are defined in the *PCI Express Base Specification*.

3.2.3. Crosstalk

All Subsystem designs must properly account for any crosstalk that may exist among the various differential signal pairs and other signals alike. Crosstalk is either near-end (NEXT) or far-end (FEXT). Each crosstalk component has potential impact on a design and must be planned for accordingly. Jitter budgets assigned to the Subsystems are set at an absolute minimum to preserve the maximum possible budget for the copper cable interconnect.

Crosstalk between differential pairs within the interconnect path influence and impact the data signals and any subsequent loss and jitter. The eye diagrams in Section 3.3 account for any/all crosstalk allowed.

3.2.4. Lane-to-Lane Skew

The skew at any point is measured at the zero crossings of differential voltage of the compliance pattern, while simultaneously transmitting on all Lanes (see Table 3-2). Refer to the PCI Express Base Specification for a definition of the compliance pattern.

Table 3-23-2. Allowable Interconnect Lane-to-Lane Skew

Jitter Parameters	Symbol	Values (RMS in ps)	Notes
Total Interconnect Skew	S_T	2.7 ns	1, 2, 3
Subsystem Skew	S_S	0.35 ns	4
Cable Assembly Skew	S_C	2.0 ns	5

Notes:

1. This does not include Transmitter output skew.
2. This is the total skew allowed for a link. If multiple cable assemblies are used, the flight time delay information (Table 6-3 bytes 108-109) must be read from all aggregated cable assemblies in the link and must meet the Total Interconnect Skew for the entire link.
3. The skew at any point is measured at the zero crossings of differential voltage of the compliance pattern, while simultaneously transmitting on all lanes. The compliance pattern is defined by the *PCI Express Base Specification*.
4. The limit is based on approximately two inches of stripline.
5. This limit is across all signal pairs. All individual cable assemblies must meet this requirement. When using multiple cable assemblies for a port aggregation, the total skew across all pairs in the link must meet this requirement.

3.2.5. Transmitter De-Emphasis

De-emphasis is required in the Transmitter for 2.5 GT/s and 5.0 GT/s to reduce ISI. As stated in Chapter 4 in the *PCI Express Base Specification*, for full swing mode, this is implemented as a -3.5 dB (± 0.5 dB) attenuation of all non-transition bits relative to the amplitude of the preceding transition bit when operating at 2.5 GT/s. For 5.0 GT/s, the *PCI Express Base Specification*, Revision 2.0 added an additional setting of -6.0 dB (± 0.5 dB) for full swing mode. For cabled implementation, transmitters must be configured for the -6.0 dB de-emphasis setting. While not optimal for all cable lengths, sufficient margin exists when using low-loss cables for reliable operation with the -6.0 dB setting.

For 8.0 GT/s, a more sophisticated transmitter equalization scheme is required. The *PCI Express Base Specification*, Revision 3.0 describes the equalization requirements, which includes not only more complex transmitter equalization, but adds receiver equalization as well. This specification builds upon these characteristics to implement this cabled form factor.

3.2.6. Skew Within the Differential Pair (Intra-Pair Skew)

Skew within a differential pair (intra-pair skew) gives rise to a common-mode signal component, which increases Electromagnetic Interference (EMI). The differential pair(s) on a Subsystem printed circuit board should be routed such that the skew within each differential pair is ≤ 0.005 inch.

Intra-pair skew of the cable assembly is more difficult to control, and tight specifications result in increased cost. A maximum skew of 0.2 UI is recommended for the cable assembly intended to operate at 2.5 GT/s. It is left up to the application to make appropriate cost/performance tradeoffs. No specific budgets are provided for intra-pair skew of the cable assembly as these are incorporated within the overall cable assembly budgets.

For 5.0 GT/s and higher, the skew is not explicitly defined for the cable. Instead, a limit is placed on mode conversion for the cable assembly in Section 6.2.4.

3.3. Eye Diagrams

The eye diagrams defined in this section represent the compliance eye diagrams that must be met for both Upstream and Downstream Subsystems. Transceiver silicon requirements are as specified in the *PCI Express Base Specification*.

3.3.1. Fixed-Port Transmit Compliance Eye

The Transmitter compliance eye is defined by the values in Table 3-3 for 2.5 GT/s, Table 3-4 for 5.0 GT/s, and Table 3-5 for 8.0 GT/s. For 8.0 GT/s, the Subsystem shall pass the eye diagram requirements with at least one of the TX equalization presets defined in the *PCI Express Base Specification*. Figure 3-2 shows the transmitter compliance eye diagram

Table 3-33-3. Transmitter Path Compliance Eye Requirements at 2.5 GT/s

Parameters	Values	Notes
Vtx _A	TBD469 mV	1, 2, 3, 5, 6, 7
Vtx _{A_d}	TBD469 mV	1, 2, 3, 5, 6, 7
Ttx _A @ BER 10-12	TBD272 ps	1, 2, 3, 4, 5
Ttx _A @ 106 Samples	TBD311 ps	1, 2, 3, 4, 5

Notes:

1. The signal acquired from TP2.
2. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating this eye diagram (see Figure 3-2).
3. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level (Vtx_{A_d}).
4. Ttx_A is the eye width.
5. The values in this table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
6. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.
7. Transition bits measured with transmitter set to -3.5 dB of de-emphasis.

Table 3-43-4. Transmitter Path Compliance Eye Requirements at 5.0 GT/s

Parameters	Values	Notes
Vtx _A	TBD 268 mV	1, 2, 3, 4, 5, 6, 7
Vtx _{A_d}	TBD 268 mV	1, 2, 3, 4, 5, 6, 7
Ttx _A @ BER 10 ⁻¹²	TBD 90 ps	1, 2, 3, 4, 5
Ttx _A @ 10 ⁶ Samples	TBD 110 ps	1, 2, 3, 4, 5

Notes:

1. The signal acquired from TP2.
2. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating this eye diagram (see Figure 3-2).
3. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level (Vtx_{A_d}).
4. Ttx_A is the eye width.
5. The values in this table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
6. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.
7. Transition bits measured with transmitter set to ~~-3.56.0~~ dB of de-emphasis.

Table 3-53-5. Transmitter Path Compliance Eye Requirements at 8.0 GT/s

Parameters	Values	Notes
Vtx _A	>34 mV	1, 2, 43 , 56 , 7
Vtx_{A_d}	TBD	1, 2, 4, 5
Ttx _A @ BER 10 ⁻¹²	>41.2 ps	1, 2, 3, 4, 5 , 7
Ttx_A @ 10⁶ Samples	TBD	1, 2, 3, 4

Notes:

1. The signal acquired from TP2.
2. The 8.0 GT/s eye requirements do not distinguish between transition and non-transition bits.
- ~~4.3.~~ An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating this eye diagram (see Figure 3-2).
- ~~2.~~ Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level (Vtx_{A_d}).
- ~~3.4.~~ Ttx_A is the eye width.
- ~~4.5.~~ The values in this table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
6. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.
- ~~5.7.~~ A worst-case reference clock with 1 ps RMS jitter is assumed for this revision of the specification. All Links are assumed active while generating this eye diagram. The eye diagram requires that the compliance pattern in 128b/130b 10 (*PCI Express Base Specification*, Revision 3.0, Section 4.2.10) is being transmitted during the test.




Figure 3-23-2. Transmitter Compliance Eye

3.3.2. Cable Port Transmit Compliance Eye

The minimum sensitivity for the Subsystem's Receiver path compliance is defined by the values in Table 3-6 for 2.5 GT/s, Table 3-7 for 5.0 GT/s, and Table 3-8 for 8.0 GT/s. The minimum sensitivity for the Subsystem's Receiver path compliance is specified in **Error! Reference source not found.** for 2.5 GT/s, Table 3-7 for 5.0 GT/s, and Table 3-8 for 8.0 GT/s. A representative eye diagram is shown in Figure 3 3.

Table 3-63-6. Receiver Path Compliance Eye Requirements at 2.5 GT/s

Parameters	Values	Notes
VrxA	TBD352 mV	1, 2, 3, 4, 6, 7
VrxA_d	TBD352 mV	1, 2, 3, 4, 6, 7
TrxA @ BER 10 ⁻¹²	TBD289 ps	1, 2, 3, 4, 5
TrxA @ 10 ⁶ Samples	TBD298 ps	1, 2, 3, 4, 5

Notes:

1. Requirements for the signal acquired from TP4.
2. All signal and timing values are referenced at the Subsystem connector mounting pads.
3. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating the eye diagram shown in Figure 3-3.
4. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level ($V_{\text{rxA_d}}$).
5. T_{rxA} is the eye width.
6. The values in this table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
7. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.

717

Table 3-73-7. Receiver Path Compliance Eye Requirements at 5.0 GT/s

Parameters	Values	Notes
V _{RxA}	TBD 211 mV	1, 2, 3, 4, 6, 7
V _{RxA_d}	TBD 211 mV	1, 2, 3, 4, 6, 7
T _{RxA} @ BER 10 ⁻¹²	TBD 82 ps	1, 2, 3, 4, 5
T _{RxA} @ 10 ⁶ Samples	TBD 102 ps	1, 2, 3, 4, 5

Notes:

1. Requirements for the signal acquired from TP4.
2. All signal and timing values are referenced at the Subsystem connector mounting pads.
3. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating the eye diagram shown in Figure 3-3.
4. Transition and non-Transition bits must be distinguished to measure compliance against the de-emphasized voltage level (V_{RxA_d}).
5. T_{RxA} is the eye width.
6. The values in this table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.
7. Transition bits measured with transmitter set to -3.5 dB of de-emphasis.

Table 3-83-8. Receiver Path Compliance Eye Requirements at 8.0 GT/s

Parameters	Values	Notes
V _{RxA}	> 34 mV	1, 2, 3, 4, 6, 7
V_{RxA_d}	TBD	1, 2, 3, 4, 6, 7
T _{RxA} @ BER 10 ⁻¹²	> 41.2 ps	1, 2, 3, 4, 5
T_{RxA} @ 10⁶ Samples	TBD	1, 2, 3, 4, 5

Notes:

1. Requirements for the signal acquired from TP4.
2. All signal and timing values are referenced at the Subsystem connector mounting pads.
3. An ideal reference clock without jitter is assumed for this specification. All Links are assumed active while generating the eye diagram shown in Figure 3-3.
4. The 8.0 GT/s eye requirements do not distinguish between transition and non-transition bits. Transition and non-Transition bits must be distinguished in order to measure compliance against the de-emphasized voltage level (V_{RxA_d}).
5. T_{RxA} is the eye width.
6. The values in this table are referenced to an ideal 100 Ω differential load, at the end of the interconnect path being tested, at the connector boundary on the Subsystem. This ideal 100 Ω differential load is implemented as two 50 Ω resistors to ground. The eye diagram is defined and centered with respect to the jitter median.
7. Maximum differential output voltage is 1.2 V as specified by the *PCI Express Base Specification*.




Figure 3-33-3. Receiver Compliance Eye

3.4. Receiver Path Sensitivity Testing

The minimum sensitivity values for the Subsystem's Receiver path compliance at 8.0 GT/s are defined in Table 3-9 and Table 3-10. The receiver path must be tested with a worst-case eye to verify that it achieves a BER $<10^{-12}$. This worst-case eye is calibrated using TX equalization settings that are optimal with the reference equalizer for the calibration channel. After calibration, the test-generator's equalization settings may be adjusted using the transmitter equalization setting in the required TX equalization space preferred by the device under test, without changing any other parameter of the test signal or recalibrating the test signal.



Note: If the test generator's TX equalization settings are adjusted away from the optimal settings and the test generator is not able to change transmitter equalization without impacting other calibrated parameters, then the other parameters must be adjusted back to the specified values.

If the test is not run in a way that produces the worst-case cross-talk that would be present with all lanes active, the additional cross-talk must be accounted for in some other way. While the receiver's capacity to adapt its own equalization is part of the test described above, its ability to request the link partner's transmitter to change its equalization settings is tested by applying a signal whose equalization settings are sub-optimal compared to the jitter sensitivity test signal described above. For this signal, the

reference receiver would not be able to achieve proper equalization by means of its own CTLE and DFE alone. Such a signal is defined using the signal resulting from the calibration method described above and adjusting the test-generator equalization. Note that if the RX under test is more capable than the reference (CTLE+DFE) receiver, the RX may not require the TX to change its equalization levels to achieve a $BER < 10^{-12}$. In any case, equalization settings resulting from this procedure ~~must be~~ used for the above RX test and, if the RX requires the TX equalization to change, such a change must be accommodated by the test set-up used. A specific methodology for this procedure is outside the scope of this specification.

733

Table 3-93-9. Free-Port Receiver Path Sensitivity Requirements at 8.0 GT/s

Parameter	Min	Max	Unit	Comments
VRX-EH-8G Eye height	> 34		mV	
TRX-EH-8G Eye width	> 0.33		UI	
Rj (Random Jitter)	3		ps RMS	
Sj (Sinusoidal Jitter) 100 MHz	12.5		ps	
Differential Mode Sinusoidal Interference 2.1 GHz	148		psmV	
Note: These are the requirements for the signal inserted into TP1.				

Table 3-103-10. Fixed-port Receiver Path Sensitivity Requirements at 8.0 GT/s

Parameter	Min	Max	Unit	Comments
VRX-EH-8G Eye height	> 34		mV	
TRX-EH-8G Eye width	> 0.33		UI	
Rj (Random Jitter)	3		ps RMS	
Sj (Sinusoidal Jitter) 100 MHz	12.5		ps	
Differential Mode Sinusoidal Interference 2.1 GHz	148		psmV	
Note: These are the requirements for the signal inserted into TP3.				

3.5. ESD

PCI Express cable ports of all link widths must withstand 2 kV of ESD contact discharge to the connector cage using the human body model (HBM), Class 2 per *JEDEC JESD 22-A114:B 2000* with power applied, without damage not limited to latch up, without a cable inserted, and without non-recoverable errors with a cable inserted.

A recoverable error is one that does not require reset or replacement of the device.

4. Interoperability

Expanding a system utilizing PCI Express over a cable requires some understanding of the implications to interoperability. Solutions with cables that are external to the Root Complex Subsystem's software or power control present some unique challenges to the PCI Express environment. A new management interface has been developed to comprehend these challenges. Flexibility in the cable assembly usage model allows users some choice in cable features

The Cable Management Interface across the cable handles sideband communication and enables power sequencing, configuration, reading of status, and an optional communications link in tandem to the PCI Express link. It is an optional feature that allows the use of cables that adhere to SFF-8449 specification to be used in an externally cabled PCI Express environment with a reduced feature set. Active Optical Cable assemblies may not implement CMI across the cable for cost or complexity reasons, and therefore optical cables may have a reduced feature set. The Upstream Subsystems should be designed in such a way as to not require the use of CMI across the cable. However, the CMI controller is still required by both the Upstream and Downstream fixed ends to read cable assembly information for configuration of the PCI Express devices that are part of the cabled link. Controlled power sequencing is only achieved with sideband enabled cable assemblies. Some features are only enabled using the sideband interface, such as SRIS or remote power control.

The Legacy adapter cable has a connector with a memory device for the Upstream end and a Legacy connector for the Downstream end. Since this version of the specification does not have all the same sideband mechanisms that the previous revisions contained, several of these sidebands must be recreated within the cable assembly.

The usage of Extension Devices (e.g., Retimers) is beyond the scope of this specification.



Note: Some Extension Devices (e.g., Retimers) are limited by the *PCI Express Base Specification* to two per link. This class of Extension Device usage is not bound by form factor specifications. This allows Extension Devices to bridge form factors, such as an expansion card interface to a cabled interface and a potential unknown number of Extension Devices in a connection hierarchy.

4.1. Cable Management Interface

Sideband signals are sent to the link partner using the Cable Management Interface. After the initial link configuration information has been obtained from the cable assembly, sideband messages may be sent from one link partner to the other if the cable supports it. A message should be sent anytime a change occurs at the local Fixed-Side. The intent is to indicate a change in status or that an action needs to be taken by the far end Fixed-Side. Sideband messages only need to be sent and monitored over the CMI for the lowest lane connector in a link that is greater than x4. The lowest lane is defined by the pin assignment in this specification and is not affected by lane reversal.

The data contained in the Vendor Message Space is beyond the scope of this specification.

The Upstream fixed device address target is address A6h.

The Downstream fixed device address target is address A4h.

4.1.1. Fixed-Side Memory Map

The fixed Subsystems have memory to provide information for configuration, status, and Vendor-Specific Messages. The Read or Write status is relative to the cable link and does not describe how it is programmed or read by the CMI controller. Table 4-1 lists the Fixed-Side memory map values.

Table 4-14-1. Fixed-Side Memory Map

Offset	Description	Values	R/W	Notes
0-2	Sideband Message Register	Bit 0 – CMI_RESET Bit 1 – CMI_POWER_ENABLE Bit 2 – CMI_SSTART Bit 3 – CMI_CLP_READY Bit 4 – CMI_CLP_PRESENT Bit 5 – CMI_WAKE Bit 6 – CMI_HOT_PLUG_ATTENTION BUTTON_DETECT Bit 7 – CMI_HOT_PLUG_INDICATOR_REQ Bits 8:23 – Reserved	<u>W</u>	The Cable Link Partner management controller writes sideband messages to these bits to indicate a status change of a sideband or indicate an action. <u>In addition, the values transmitted in a sideband message are stored in the CPSR for the Cabled Link Partner to read when the CMI is free. A sideband message is retransmitted by a CMI controller at a defined interval regardless of a change in any status bit within the CPSR and Sideband Message.</u>
3	Sideband Message CRC	Bit 0 – CPSR_RESET Bits 1:2 – Reserved Bit 3 – CPSR_CLP_READY Bit 4 – CPSR_CLP_PRESENT Bits 5:15 – Reserved 8 bit CRC	<u>RW</u>	See Section 4.1.2.9
<u>0-2</u>	<u>CPSR Register</u>	Bit 0 – CPSR_RESET Bit 1 – CPSR_POWER_ENABLE Bit 2 – CPSR_SSTART Bit 3 – CPSR_CLP_READY Bit 4 – CPSR_CLP_PRESENT Bit 5 – CPSR_WAKE Bit 6 – CPSR_HOT_PLUG_ATTENTION BUTTON_DETECT Bit 7 – CPSR_HOT_PLUG_INDICATOR_REQ Bits 8:23 – Reserved	<u>R</u>	<u>This register stores the Cable port status.</u>
<u>3</u>	<u>CPSR_CRC</u>	<u>8 bit CRC</u>	<u>R</u>	See Section <u>4.1.2.11</u>

Offset	Description	Values	R/W	Notes
4-5	PCI Express Cable Port Capabilities Register	Bit 0 – 5.0 GT/s Data Rate Supported Bit 1 – 8.0 GT/s Data Rate Supported Bits 2:4 – Reserved Bit 5 – SRIS Support Bits 6:15 – Reserved	R	These bits are controlled by the Fixed-Side and used by the opposite side during configuration.
6-7	Vendor ID	PCI SIG defined Vendor ID in hex	R	Use for identifying common Boxes for port configuration. Recommended to match System Vendor ID.
8-9	Device ID	Vendor defined Device ID in hex	R	Use for identifying common Boxes for port configuration. Recommended to match Subsystem ID.
10	Connector Configuration	Bits 0:1 – Physical port ordering of x4 connectors in a larger link. Encoding: – 00b = Lanes 0-3; – 01b = Lanes 4-7; – 10b = Lanes 8-11 – 11b = Lanes 12-15. Bits 2:7 – Value used to identify connector ports that are permitted to be combined into a single link.	R	Lower two bits indicate physical port ordering of x4 ports that may be used to create a larger link, up to x16 lanes or four physical ports. Upper six bits indicate physical ports that may belong to a single link. For Upstream Fixed-Side devices, physical ports are allowed to bifurcate into a maximum of four individual link connections. Physical ports with matching values may be combined. The specific values are implementation-specific.
11-18	Serial Number	Vendor defined 64-bit field to uniquely identify each unit of a model.	R	Uniquely identify vendor Boxes of same Vendor ID / Device ID. Used for port configuration connectivity (plugging cables incorrectly when employing cable aggregation).
19-42	Vendor ID Text	Vendor-defined 24-character string value in ASCII.	R	
43-66	Vendor Description Text	Vendor-defined 24-character string value in ASCII.	R	
67-126	Reserved		R	
127	CRC		R	CRC covers static data from Byte 4 to Byte 126.
128-255	Vendor Message Space	Vendor Specific	-R/W	This space is used for vendor-specific communication.



Implementation Note: The data in the Vendor Message Space may include pointers to additional memory that a vendor may wish to take advantage of. Usage of this space is beyond the scope of this specification.

4.1.2. Sideband Messages

Sideband messages are used as an economical alternative to connector pins and cable wires. This method enables the transmission and reception of sideband messages to update a status, indicate an action, or inform the CLP of a local status change. A message must be sent by the CMI controller when a status change occurs in the Fixed-Side Subsystem that is to be reflected via a sideband. If transmission of the message fails, the CMI controller must immediately attempt to resend the sideband message until the transmission is successful, or until MGTPWR is removed.

When sideband messages are sent by the CMI controller, any bits not specified to change must reflect their state from their last transmission that specified their change or if they were never changed, from initialization.

Messages are sent periodically even when no bits have changed. If no bits have changed, the previous message is repeated. A message should always be sent within $T_{SB_TIMEOUT}$ of the previous message to ensure a cable disconnect or CLP power loss is detected.

Table 4-24-2. Sideband Messages

Sideband Message Bits	Bit	Values	R/W	Notes
CMI_RESET	0	Set indicates PCI Express Reset is asserted, Cleared indicates PCI Express Reset is de-asserted.	W	This signal indicates the state of the reset in the Root Complex Box and is sent by the Upstream Fixed-Side only. If sent by the Downstream Fixed-Side, <u>it</u> is undefined.
CMI_POWER_ENABLE	1	Set indicates that main power should be applied, Cleared indicates main power should be removed.	W	Indicates the directed power state and is sent by the Upstream Fixed-Side only. If sent by <u>the</u> Downstream Fixed-Side, <u>it</u> is undefined.
CMI_SSTART	2	Set indicates a Box should wait to enable main power until a sideband message is received with CMI_CLP_READY set. Cleared indicates that main power may be enabled as soon as directed.	W	This bit allows staggered power up of Boxes that are logically connected via PCI Express, but do not share a common power delivery or control structure. This is sent by the Upstream Fixed-Side only. If sent by <u>the</u> Downstream Fixed-Side, <u>it</u> is undefined.
CMI_CLP_READY	3	Set indicates that Cable Link Partner is ready for communication (e.g., it has main power enabled, and clock is stable).	W	Indicates that the source is ready for communication. This is a status signal and does not indicate action, except when CMI_SSTART is Set. This may be sent by either end of the cabled link.
CMI_CLP_PRESENT	4	Set indicates local Box is present, Cleared indicates local Box is absent.	W	This indicates to the CLP that the local Box is present.

Sideband Message Bits	Bit	Values	R/W	Notes
CMI_WAKE	5	Set to request that Upstream Fixed-Side enables main power.	W	Driven by the Downstream Fixed-Side only. If sent by <u>the</u> Upstream device the result is undefined. See Table Note.
CMI_HOT_PLUG_ATTENTION_BUTTON	6	Cleared- No event to report Set - Hot-Plug Attention button event has occurred.	W	This bit indicates that the opposite Fixed-Side has detected a button press.
CMI_HOT_PLUG_INDICATOR_REQ	7	Cleared: Normal Operation Set: Hot-Plug Indicator Request	W	This bit indicates that the opposite Fixed-Side has received direction to blink the CLP's indicator.
Reserved	8-15			

Note: Due to the slow nature of the Cable Management Interface, OBFF is not implemented as a sideband message. This is in contrast to other form factors that implement OBFF via the #WAKE sideband signal. OBFF messages sent in-band are not impacted by this form factor specification.

4.1.2.1. CMI_RESET

The CMI_RESET bit indicates the status of the PCI Express reset of the Upstream Fixed-Side Subsystem. The bit is Set when local reset is asserted and Cleared when de-asserted. The CMI_RESET bit is sent by the Upstream Fixed-Side only. All DSF CMI controllers must send messages with CMI_RESET Set from controller initialization until a message is received on the USF port with it cleared. When sent by the Downstream Fixed-Side Subsystem, the meaning is undefined; therefore, the Upstream Fixed-Side takes no action based on this. ~~the meaning is undefined.~~

4.1.2.2. CMI_POWER_ENABLE

The CMI_POWER_ENABLE bit controls the power-up of Boxes that are logically connected via PCI Express but do not share a common power delivery or control structure. This bit is sent by the Upstream Fixed-Side only. When sent by Downstream Fixed-Side Subsystem, the meaning is undefined; therefore, the Upstream Fixed-Side takes no action based on this. ~~is undefined.~~



Implementation Note: Boxes with only DSF ports send this as a reflection of a local POWER_GOOD signal. However, if this simply reflects POWER_GOOD then a predetermined power up sequence is not enabled and one or more Boxes in the cabled hierarchy may not be ready to respond to a valid PCI Express Configuration Request with a Successful Completion.

4.1.2.3. CMI_SSTART

CMI_SSTART allows for a staggered start of Boxes connected via PCI Express cables to both reduce power loading at initialization and controlled power sequencing. The CMI_SSTART bit controls whether the action of CMI_POWER_ENABLE or CMI_WAKE is delayed until CMI_CLP_READY is set. Setting the CMI_SSTART bit in the message only has meaning when either CMI_POWER_ENABLE or CMI_WAKE is also being Set. If Set in a message without a transition of either CMI_POWER_ENABLE or CMI_WAKE being Set, the meaning of the CMI_SSTART bit is undefined and no staggered start will be initiated. The bit can be Cleared any time after a message has been received with CMI_RESET de-asserted.

4.1.2.4. CMI_CLP_READY

CMI_CLP_READY is sent by a CMI controller to the CLP to indicate the local Box is ready to receive a message with CMI_RESET de-asserted. This implies that local power is stable and within operating parameters, clocks are stable, the configuration data from the cable assembly has been obtained, a message with CMI_CLP_READY has been received on all DSF ports that can receive sideband messages, and any other implementation-specific criteria has been met for the Box to begin receiving PCI Express transactions.

The CMI_CLP_READY signal must remain Set until the unit is no longer able to maintain an operational PCIe interface. PERST# could cause the conditions listed here but should not cause READY to go away. It is permitted for the CMI controller to Clear CMI_CLP_READY for implementation-specific reasons, such as loss of key functionality. The behavior of the other Boxes upon receipt of a message with CMI_CLP_READY Cleared after power-up sequencing is implementation-specific.

4.1.2.5. CMI_WAKE

This bit indicates that the Downstream Fixed-Side is requesting the reactivation of main power. The Upstream Fixed-Side Subsystem must assert its local WAKE# signal and propagate (if possible) the wake message Upstream. CMI_WAKE bit Set is accompanied by CMI_SSTART bit Set to indicate a staggered power-on behavior. CMI_WAKE is driven by the Downstream Fixed-Side System only. CMI_WAKE is undefined when sent by an Upstream device and no power on action will be taken ~~is undefined when sent by an Upstream device.~~ See Section [4.24.1.3](#) for details of power sequencing.

4.1.2.6. CMI_CLP_PRESENT

This bit is used to indicate the presence of the local Box and is always Set when sideband messages are sent. Use of this bit for other behaviors is beyond the scope of this specification.

4.1.2.7. CMI_HOT-PLUG_ATTENTION_BUTTON

This bit indicates that the CLP has detected an Attention Button press. Refer to Section 6 of the *PCI Express Base Specification* for more information

4.1.2.8. CMI_HOT-PLUG_INDICATOR_REQ

This bit indicates that the CLP has received a request to blink its indicator.

4.1.2.9. Sideband Message CRC

The Sideband Message CRC is an 8-bit cyclic redundancy check (CRC-8) of each write transaction. The CRC is permitted to be calculated in any way that adheres to a CRC-8 represented by the polynomial, $C(x)=x^8 + x^2 + x^1 + 1$, and must be calculated in the order the bits are received. All data between the first Start and the Stop bits, excluding the ACK, NACK, and Repeated Start bits, are included in the CRC. This includes the memory device address, memory location within the device, and Write bit.

4.1.2.9.4.1.2.10. Cable Port Status Register

The Cable Port Status register (CPSR) indicates status of various states of the port. Some of these bits reflect the bits sent in a sideband message to the CLP. It is strongly recommended that the CPSR be updated when the sideband message is first transmitted, even if the transmission is retried. If the sideband message transmission is delayed for any reason, such as the Box losing bus arbitration, the CPSR is still ~~be~~ updated.

4.1.2.11. Cable Port Status Register CRC

The Cable Port Status Register CRC is an 8-bit cyclic redundancy check (CRC-8) of each write transaction. The CRC is permitted to be calculated in any way that adheres to a CRC-8 represented by the polynomial, $C(x)=x^8 + x^2 + x^1 + 1$, and must be calculated in the order the bits are received. The memory device address, memory location within the device, and Read bit shall be included in the CRC. See Section 4.1.2.9 for more clarification.

4.1.2.10.4.1.2.12. PCI Express Cable Port Capabilities Register

Each Fixed-Side Subsystem contains a register space containing configuration information that is used to configure the system (see Table 4-3). The CLP may use this information to configure its PCI Express settings.

Table 4-34-3. Fixed-Side Capability Registers

Status Bit	Location	Description	Notes
5.0 GT/s Support	0	Support for 5.0 GT/s operation	5.0 GT/s operation is optional. Must be Set to 1b if 5.0 GT/s is supported.
8.0 GT/s Support	1	Support for 8.0 GT/s operation	8.0 GT/s operation is optional. Must be Set to 1b if 8.0 GT/s is supported.
Reserved	2		
Reserved	3		
Reserved	4		

Status Bit	Location	Description	Notes
SRIS Support	5	Support for Separate Reference Clocks with Independent Spread Spectrum	This is an optional feature. Both sides must indicate support for this in the PCI Express Cable Port Capabilities register.
Reserved	6 to 15		

4.1.3.4.1.2.13. Static Data CRC

The Static Data CRC is an 8-bit cyclic redundancy check (CRC-8) of each read transaction. The CRC is permitted to be calculated in any way that adheres to a CRC-8 represented by the polynomial, $C(x)=x^8 + x^2 + x^1 + 1$, and must be calculated in the order the bits are received. Address and command bits shall be included in the CRC. The address is based on the value of CADDR (see [Section 4.1.2.11 for more details about CRC implementation](#)).

4.2. Power Sequencing

Power sequencing for cabled PCI Express applications must comprehend multiple Boxes with independent power supplies, resets, and firmware. Provisions are made to work with host systems that do not enable true power sequencing of the entire cabled system.

4.2.1. Power-Up Sequencing

CPSR_RESET must be Set by the management device when MGTPWR is applied. Downstream Subsystems should be powered before the Upstream Subsystem to allow the Downstream Subsystem time to fully initialize before the PCI Express link begins the discovery process.

The power-up sequence for the cable link and PCI Express devices must follow the following sequence:

1. Enable MGTPWR; there is no precondition for enabling MGTPWR.
2. Initialize CMI Controller.
3. Set CPSR_CLP_PRESENT in the ~~Fixed-Side-Status~~Cable Port Status register. Set CPSR_RESET for DSF Port(s).
If the subsystem has a USF port supporting sideband messages, local platform reset of PCI Express devices must reflect the state of the sideband message CMI_RESET received, if any, on the USF Port; if sideband messages are supported; otherwise the *PCI Express Base Specification* is followed.
4. Wait to proceed until assertion of CBLPRSNT#.
5. After a maximum delay of TMGT_INIT from the assertion of CBLPRSNT#, as detected by the CMI Controller, read Cable Memory.
 - Upstream Subsystem may attempt to read configuration data from address A2h or Address A0h. If there is no response from address A2h, the Upstream Subsystem must read from address A0h.
 - Downstream Subsystems must read from address A0h.

See Section 2.6 for more details regarding the configuration information. The CMI Controller is permitted to perform checks on the integrity of the data before proceeding to the next step.



Observation: The reading of the cable memory is necessary for proper configuration of the PCI Express device that interfaces to the cable port. However, there may be one or more failed attempts to read the data due to bus contention or non-response. The read process must be attempted until the configuration data is successfully acquired.

6. If sideband messages are supported (based on the support for Sideband Messages bit in the cable assembly's memory), read CLP Fixed-Side Memory. If the cable supports sideband messages, as indicated by the support for Sideband Messages bit in the Cabled PCI Express Capabilities 1 register, the management controller is permitted to attempt to detect the presence of the CLP by reading the opposite Fixed-Side via CMI. This acts as a Hot-Plug detection scheme and is located at the same place in the memory map as the target for the sideband message. A CMI controller updates the CPSR and then transmits the same data in a sideband message to the CLP. The CLP may read the CPSR anytime CMI is idle.



Note: Bytes 4-18 (PCIe Port Capabilities Register, Vendor ID, Device ID, Connector Configuration, and Serial Number) must be processed to correctly establish the system configuration.

7. Prepare to Configure Core Clocking components and PCI Express devices.
- a) The default behavior of Fixed-Side Subsystems is to:
- Support SRNS. If the Fixed-Side cable port supports SRIS and the PCI Express Cable Port Capabilities register Bit 5 of the CLP indicates support of SRIS then the CMI controller must change the local reference clock to operate in SRIS mode. The port must deassert CMI_CLP_Ready, configure the clock, (refer to the PCI Express Base Specification), and then and then reassert CMI_CLP_Ready.
 - All physical cable port link widths and locations are set to a default configuration that is implementation-specific.
 - Cable link widths may link train to a smaller link width (e.g. a x8 link may train to be a x4 link), via normal PCI Express Link Training protocols, if the cable assembly link width is smaller than the cable port link width. Lane 0 for any link must be connected to an effective Lane 0 of the CLP. Support of cable aggregation is not required.
- b) Port Width Configuration
- Bifurcation of the Upstream Port

CMI Controllers for DSF Ports that support cable bifurcation must obtain information that allows for port configuration from the CLP and the cable assembly before the PCI Express link initialization begins (refer to the PCI Express Base Specification).

 - > DSF logical ports are permitted to be mapped to cabled ports in which the following fields of the CLP for each physical cable port match: Vendor ID, Device ID, Serial Number, and the Connector Configuration bits 2-7.
 - > Bifurcation of a port, beyond an implementation's default configuration, with a cable that does not support sideband messages is beyond the scope of this specification.
 - > Bifurcation of the USF Port is beyond the scope of this specification.

— Cable Configuration

Cabled links are permitted to aggregate multiple cable assemblies to form a single link. CMI Controllers for Ports supporting cable aggregation must be able to obtain information that allows for port configuration from the CLP and the cable assembly before the PCI Express hardware initialization begins. Cable aggregation with cable assemblies that do not support sideband messages is beyond the scope of this specification. It is implied that cables supporting sideband messages are PCI-Express specific and have Byte 111, bit 0 set, but this may be checked as a first pass for configuration.

Aggregation of Active Cable Assemblies is beyond the scope of this specification. The CMI controller compares the cable information across all physical ports connected to a logical port to determine if aggregation of the inserted cable assemblies is permitted. If the cable assembly supports the maximum link width (based on Byte 112 of the cable assembly's memory, Cabled PCI Express Capabilities 2 register bits 0:2, of the logical port), further checks are not required. If any cable assembly inserted to the logical port is smaller than the logical port width, beginning with lowest lanes of the logical port, cables must contain matching data in the following bytes of the cable memory map, unless otherwise specified:

≥ —Base and Span from the Propagation Delay fields

≥ —Cable Technology Type field

≥ —Attenuation fields

— If the cables inserted to consecutive sets of x4 lanes of a logical cable port, identified by the port order of the CLP, do not meet matching criteria, then only the lanes connected to consecutive cable assemblies that meet the matching criteria, starting with the lowest lanes, are permitted to be enabled.

> It is recommended the CMI Controller use bits 0:1 of the Connector Configuration register to verify aggregated cables are connected in the correct order such that the lanes of the Link are correctly ordered from 0-n to the other side.

> It is recommended that the PCI Express device receivers of the disabled lanes be ignored to prevent unknown states or configurations within the PCI Express Device.

> Determine Equalizer presets and hints based on Section 4.2.1.1. It is permitted for the presets and hints to be programmed at this stage, if the implementation allows.

It is recommended to also change any settings that may be impacted by cable delay for both PCI Express devices and the Root Complex. Examples of these settings may include, but are not limited to, replay timers, flow control credits, and presence detect circuit timings.

8. Enable Power Main.

Main power for any Box with a Fixed-Side Subsystem may be enabled either via a sideband message or local power control. It is recommended that Boxes with local power control support initiation of staggered power sequencing over the cabled links to ensure all Downstream devices are ready for access when the Box with the Root Complex de-asserts Reset and to minimize potential strain on external power systems during the power-up process. The specific mechanism to enable this feature when using local power control is beyond the scope of this specification.

The CMI_SSTART bit in the sideband message enables the staggered start feature. This enables sequential power enablement of Box in the cabled hierarchy. All Boxes with a cabled port must support the propagation of CMI_SSTART via sideband messaging. All Boxes with a USB port must support the staggered start feature when directed via sideband message. Implementing staggered start via local control is strongly recommended.

There are two staggered start implementation options: DSF Gated and Local First. The choice of which implementation to support or use is implementation-specific. The DSF Gated implementation must send sideband messages with CMI_RESET Set, CMI_POWER_ENABLE Set, and CMI_SSTART Set from all DSF Ports, and then receive a sideband message with CMI_CLP_READY Set from all those DSF Ports before main power is permitted to be enabled. The Local First implementation permits local power to be enabled before sending sideband messages with CMI_RESET Set, CMI_POWER_ENABLE Set, and CMI_SSTART Set to all DSF Ports. With either method, all DSF with CBLPRSNT# asserted, on a Subsystem with a USB PORT ports on Subsystems with USB ports MUST receive messages with CMI_CLP_READY Set ~~on all the USB ports~~ before de-asserting local resets. The choice of which implementation to support or use is implementation-specific.



Implementation Note: To enable further power staggering and reduce initial current draw, Boxes with multiple DSF ports may stagger sending the remote power messages across its DSF ports that are not ready.

Whether or not CBLPRESENT# is asserted, sideband messages are not required to be sent to any port with a cable assembly that does not support sideband messages. ~~"Sideband messages are not required to be sent to ports that show CBLPRESENT# asserted but the cable assembly does not support sideband messages.~~



Observation: If the Root Complex does not receive a Successful Completion to a valid Configuration Request within 1.0 s of the Root Complex exiting Conventional Reset, the Root Complex is permitted to determine the Downstream Box's PCI Express device is broken. Usage of sideband messages enables power sequencing of the Boxes, such that all Downstream Box's cabled ports are ready for configuration before the Root Complex begins sending Configuration Requests to cabled ports. Usage of cables that do not support sideband messages do not enable specific sequencing; Boxes, such that power up in an undetermined sequence, and some or all Downstream Boxes may not be ready to meet time requirement for a Successful Completion.

1013

Main power is permitted to be enabled for:

- Boxes with USF Port only
 - Enable main power when the Box is directed to do so, ~~by~~ either by local control or if the USF port receives a sideband message with CMI_POWER_ENABLE Set.
- Boxes with DSF Port(s) only
 - If the Box receives a message with CMI_WAKE Set.
 - If CMI_SSTART is also Set, the CMI controller implements the DSF Gated implementation of staggered start.
 - The Local First implementation is permitted if the CMI controller prevents platform reset de-assertion until ~~the Downstream Boxes are ready~~ a message with CMI_CLP_READY Set is received on all DSF ports with CBLPRSNT# Set.
 - If directed by local power control to enable main power,
 - If either the cable inserted to the DSF port does not support sideband messages or the staggered start feature is not to be initiated, then main power is permitted to be enabled immediately.
 - If the staggered start feature is to be initiated, the CMI controller implements the DSF Gated implementation of staggered start. The Local First implementation is permitted if the CMI controller prevents platform reset de-assertion until ~~the Downstream Boxes are ready~~ a message with CMI_CLP_READY Set is received on all DSF ports with CBLPRSNT# Set.



Implementation Note: Cabled ports provided via an adapter card, or module, may not have the ability to gate power for the Box in which the adapter is inserted. A power switch may be provided by the adapter to have the CMI controller initiate the remote power enable feature and, ~~(1 status)~~ once the DSF port(s) receive messages with CMI_CLP_READY, the local wake signal for the Box may be asserted according to form factor requirements.



Observation: The Box with only a DSF Port contains the Root Complex. If sideband messages are not supported by the cable, then Downstream Boxes may not configure properly before the Box de-asserts platform reset, unless the Downstream Boxes are powered on first by other means.

- Boxes with both a USF Port and at least one DSF Port.



Implementation Note: Either DSF Gated or Local First are permitted to be implemented by this type of Box.

- If directed by local power control
 - > If either the cable inserted to the DSF port does not support sideband messages or the staggered start feature is not to be initiated, then main power is permitted to be enabled immediately.
 - > If staggered start is to be initiated, either method is permitted.
 - ~~If directed~~ via a DSF Port by receiving a sideband message with CMI_WAKE Set
 - > If the sideband message also has CMI_SSTART Set, either staggered start implementation is permitted.
 - > If the sideband message also has CMI_SSTART Cleared, main power is permitted to be enabled, and sideband messages must be sent to all other DSF Ports in the Box with both CMI_POWER_ENABLE and CMI_RESET Set, but CMI_SSTART Cleared. The sequencing of sending this message to the DSF ports is not specified. A message with CMI_WAKE Set is permitted to be sent to the USF Port.
 - ~~If directed~~ via USF Port ~~be by~~ receiving a sideband message with CMI_POWER_ENABLE Set
 - > If the received sideband message also includes CMI_SSTART Set, either staggered start implementation is permitted.
 - > If the received sideband message includes CMI_SSTART Cleared, a sideband message with both CMI_RESET and CMI_POWER_ENABLE Set, and with CMI_SSTART Cleared, must be sent to all DSF Ports with CBLPRSNT# asserted; and is permitted to enable main power immediately.
9. Set CPSR_CLP_READY when main power is within operating limits, clocks are stable, cable memory configuration data has been obtained, and the Box is ready for CMI_RESET to be de-asserted.



Implementation Note: Implementations of a Box are permitted to have additional criteria gate CPSR_CLP_READY. These may include initialization of functions outside the PCI Express specific functions.

10. Indicate to Cable Link Partner that Box is ready.
- If power was initiated either by receiving a message with CMI_WAKE Set on a DSF Port ~~set~~, or by local power control
 - Send a message ~~to from~~ the USF port with both CMI_WAKE and CMI_CLP_READY Set. CMI_SSTART must reflect the state of CMI_SSTART from the message on the DSF port with CMI_WAKE Set, or the staggered start implementation for the local power control mechanism.
 - Send a message ~~to from~~ all DSF ports with CMI_CLP_READY Set.
 - If power was initiated by receiving a message from the USF Port with CMI_POWER_ENABLE Set, send message from all ports with CBLPRSNT# asserted and CMI_CLP_READY Set; it must reflect the state of CMI_SSTART with CMI_POWER_ENABLE Set.

1078 11. CMI_WAKE is to be Cleared from CMI messages once a message is received with CMI_RESET
1079 Cleared.

4.2.1.1. Equalizer Settings Based on Cable Memory

Cabled PCI Express ports need to be configured to work correctly in the cabled environment. Loss for the cable is stored in the cable memory map. For signal rates above 5.0 GT/s, the Fixed-Side CMI controller sets the transmitter and receiver settings for 8.0 GT/s based on loss value for 4 GHz. The recommended settings for cable port transmitter and receiver starting values are listed in Table 4-4. These values are derived from passive cable assemblies.

Table 4-4. Cable Port Transmitter and Receiver Starting Values

Cable Signal Transmission (in dB)	Receiver CTLE Settings
-3 or higher ^{1,2}	-6
-4	-7
-5	-7
-6	-8
-7	-8
-8	-9
-10	-10
-11	-11
-12 or lower ³	-12

Notes:

1. Cable assemblies reporting a transmission loss value of 0 dB are electrically buffering the signal before exiting the cable assembly and indicate a full-swing signal is being sent from the cable assembly to the Fixed-Side Subsystem. Cable assemblies are not permitted to report a transmission loss value larger than 0 dB (i.e., the signal exiting the cable assembly cannot be larger than full-swing).
2. It is recommended that inserted cable assemblies reporting a transmission loss value greater than -3 dB be set to the minimum receiver value of -6 dB. Receivers that support Rx CTLE values of greater than -6dB are beyond the scope of this specification.
3. It is recommended then that inserted cable assemblies reporting a transmission loss value less-less than -12 dB be set to the maximum receiver CTRL value of -12 dB. Receivers that support Rx CTLE values of less than -12 dB are beyond the scope of this specification.

It is recommended that transmitters begin with a preset setting of P6 unless the loss of the cable is reported as 0dB. A setting of 0dB implies a fully buffered cable and, in this case, ~~and~~ it is recommended that the transmitter's preset setting be configured for P4. The settings of P4 and P6 are defined in the PCI Express Base Specification.

Cable technology also impacts the presets. Byte 147 bits 4-7 indicate the Cable Technology Type. Values of Ah, Bh, or Fh indicate the total effective loss of the cable provided by bytes 186-189. Values of 0h or Ch indicates that the cable is fully buffered (i.e., loss is 0 db). It is recommended that ports connected to fully buffered cable assemblies set the Tx preset setting to P4 and Rx CTLE to its minimum value. A value of Dh indicates the Rx end is fully buffered (or non-linear), but the Tx end is not.

The receiver should be set to the minimum ~~CTRL-CTLE~~ setting while the transmitter should be optimized for the loss of the transmitter PCB interconnect and the reported cable loss. Refer to the *PCI Express Base Specification* for the definitions of the presets.

The method for the CMI controller to program the PCI Express component is implementation specific and beyond the scope of this specification.

4.2.1.2. PCI Express Configurations Based on Cable Memory

Long cable delays impact the link performance and the PCI Express protocol. Replay Timers and Completion Timers expire if responses are not received in an expected timeframe. Limited credit allocation may limit the amount of data in flight over long cable reaches. The Propagation Delay field provides information that is used to set timers and credits for optimal performance. It is recommended that if the propagation delay is greater than 20 ns, then, for all Devices capable of issuing requests that will traverse the cable, the Completion Timeout value field in the Device Control 2 register be set larger than 10 ms (e.g., to Range B or higher). Devices should be programmed with the maximum number of credits possible to allow for as many packets to be in flight as possible. For any type of active cable, indicated by the Cable Technology Type field, it is recommended, for ports connected to a cable, to set the Extended Synch bit to 1, in the Link Control register, and to program the initial N_FTS value to 255, to allow maximum time for link training and link recovery. It is also recommended that ASPM be disabled for all cable technologies, except passive cables, to avoid gain and noise problems on the link.

~~When using Latency Tolerance Reporting (LTR), software must consider the cable assembly's propagation delay, in programming platform latencies.~~

~~When isochronous traffic is being supported, software must consider the cable assembly's propagation delay, in programming isochronous transaction latencies.~~



Implementation Note: When using Latency Tolerance Reporting (LTR), software must consider the cable assembly's propagation delay, in programming platform latencies.

When isochronous traffic is being supported, software must consider the cable assembly's propagation delay, in programming isochronous transaction latencies.

4.2.2. Power-Down Sequencing

4.2.2. The power down sequence involves both local devices and, if sideband messages are supported, the transmission & reception of sideband messages. If sideband messages are not supported by the cable assembly, then implementation specific means will be necessary for power control in conjunction with registers and in-band messages for the PCIe links.

~~Figure 4-2 illustrates the power-down sequencing resulting from a graceful or Hot-Unplug event.~~

4.2.2.1. Sequence Steps

Following are the sequence steps for a graceful shutdown:

- ❑ PCI Express lanes transitions to inactive state (Device in D3_{hot})
- ❑ The Sideband message is initiated by the Box with the Root Complex with CMI_RESET Set. The message is propagated from all USF Ports to any DSF Ports immediately, taking priority over any other sideband messages that may be queued.
- ❑ CMI controller for a Box with a USF Port asserts local reset signal.
- ❑ The Sideband message is initiated by the Box with the Root Complex with both CMI_POWER_ENABLE and CMI_CLP_READY Cleared. The message is propagated from all USF Ports to any DSF Ports.
- ❑ CMI controller for a Box with USF Port(s) is permitted to disable main power



Observation: Most Boxes are expected to follow the status of CMI_POWER_ENABLE that originates in the Box with the Root Complex, after initial Reset de-assertion. There are implementations in which it is desirable to leave the Box powered without PCI Express communication. Boxes with USF ports are permitted to leave main power enabled when a sideband message is received with CMI_POWER_ENABLE Cleared. However, the PCI Express device must follow the state of CMI_RESET.



Note: It is recommended that the Downstream Subsystem is not powered off before the Upstream Subsystem, if using local power control, to allow system software to gracefully exit before the hardware is effectively removed from the system topology.

4.2.3. Cable Management Interface Timings

The operational description of the CMI is specified by the *PCI Express OCuLink Specification*, Appendix B. General protocol and Read/Write operations are specified. In addition to these parameters, the arbitration protocol from the I2C Revision 2.1 is followed to support the multi-master nature of the CMI and potential bus contention between the Fixed-Side CMI controllers at each end of a cabled link.

4.2.4. Link Power Management

The *PCI Express Base Specification* specifies Link Power Management states. The various form factors are permitted to define which states are supported, optional, or not supported, some of which are required by all existing form factor specifications while others may be optional. A cabled PCI

1152 Express Port, adhering to this revision of the specification, may be implemented on top of any other
1153 form factor. Support of Link Power Management states are a system-level requirement and are not
1154 imposed by this specification.

Downstream Subsystems wishing to guarantee compliancy with all possible system-level specifications must support all defined Link Power Management states.



Note: Some active cables do not respond well to lack of signal transitions during the low power states defined by the *PCI Express Base Specification*. *In such cases, systems should disable the Link Power Managing states to increase compatibility and improve performance.*

4.3. PCI Express Features

As the cable assembly in this version of the specification does not have all the same sideband mechanisms of other form-factor specifications, some PCI Express features are not supported by devices connected via a cable.

Support is not provided by this specification for the following features defined in the PCI Express Base Specification:

- ❑ No support for dynamic control of reference clock via Clock Power Management during L1 states:
 - Due to the lack of a sideband signal corresponding to “CLKREQ#” available on other form-factors:
 - Firmware must not program the USF’s Enable Clock Power Management bit to enable dynamic clock control. The Clock Power Management field should have a value of 0b.
- ❑ Limited ASPM Support
 - Due to the lack of a sideband signal corresponding to “CLKREQ#” available on other form-factors:
 - Firmware must not program the Fixed-Side cable port’s L1 PM Substates Control 1 register enable bits (i.e., any of PCI-PM L1.2 Enable, PCI-PM L1.1 Enable, ASPM L1.2 Enable, and ASPM L1.1 Enable) to enable L1 PM Substates. *However, firmware is still permitted to program the Fixed-Side cable port’s ASPM Control field to enable L1, if both the Fixed-Side cable port and the CLP indicate support for ASPM L1.*
 - Support for dynamic control of reference clock via Clock Power Management during L1 states, as *Mechanical Definition* described in the *PCI Express Base Specification*, is not provided by this specification.
 - It is recommended that all ASPM Substates be disabled via the Active State Power Management Support field in the Link Capabilities Register.
- ❑ No Support for OBFF using WAKE# signaling:
 - *This is due to the lack of a sideband signal corresponding to “WAKE#” available on other form-factors. So even though the Cable Management Interface recreates the CMI_WAKE sideband, the interface is too slow to use this for OBFF signaling. The OBFF Supported field in the Device Capabilities 2 Register must be set to either 00b or 01b.*
 - Firmware must not program the Fixed-Side cable port’s OBFF Enable field to use WAKE# signaling. *However, firmware is still permitted to program the Fixed-Side cable port’s OBFF Enable to use Message signaling, if the Fixed-Side cable port indicates support for that method.*
 - DSF ports shall set the Slot Implemented bit in the Capabilities Register to indicate to software *that this is not an integrated link.*

5. Fixed-Side Definition

This specification defines distinct connectors to support multiples of x4 Link width. Links with lower lane counts must use this connector and have two or more unused Lanes. Links greater than x4 Lanes (i.e., x8, x12, and x16 Lanes) use multiple connectors to achieve the total Link width. Cables that contain the entire link width are referred to as monolithic assemblies. An example would be a x16 cable. The x8 cable is considered monolithic ~~if used in a x8 link, if used in a x8 link; however, it is not if used as part of a x16 link. however, is not if used as part of a x16 link.~~

5.1. Fixed-Side Channel Characteristics

The cabled channel was developed using $100\ \Omega$ as a target for both PCB's and cables. The tolerance was $\pm 10\ \Omega$. The loss target is -2.3 dB at 4 GHz. Care should be taken to minimize the effect of via stubs in the channel, as these can have significant impact on channel performance.

5.2. Signal Description

The external PCI Express cable connector and cabling support the following signals:

- PETpN/PETnN (required): ~~labeled~~ PCI Express Transmitter pair(s), ~~labeled~~ where **N** is the Lane number (starting with 0); ~~“p”~~**p** is the true signal ~~while and “n”~~**n** is the complement signal.
- PERpN/PERnN (required): ~~labeled~~ PCI Express Receiver pair(s), ~~labeled~~ where **N is the Lane number (starting with 0); p is the true signal and n is the complement signal N is the Lane number (starting with 0); “p” is the true signal while “n” is the complement signal.**
- Auxiliary signals, as defined in Chapter 2, Auxiliary Signals.

The Fixed-Side PCI Express signals must be connected to the PCI Express device, transmitters to the PETxY pins and receivers to the PERxY pins, as described in Table 6-2. The cable assembly provides the cross-over path that enables transmit pins to send data to receive pins.

5.3. Fixed-Side Connector

Refer to the *SFF-8644, Mini-Multilane 12 Gbs 8/4x Shielded Connector Specification* for mechanical information.

5.3.1. Pin-out

Table 5-1 contains the pinout of the base connector. If the implementation utilizes more than a single x4 fixed connector, the pinout should follow that shown in Table 5-2. Figure 5-2 shows the four Fixed-Side x4 connectors ~~pin assignments configured~~ to enable a x16 link. Figure 5-3 illustrates the ordering of the connectors on a card edge.

The connector pin definitions are referenced to the PCI Express device pins of the same Fixed-Side Subsystem. PETxY pins of the connector must be connected to the transmitter pins of the device. PERxY pins of the connector must be connected to the receiver pins of the PCI Express device.



Note: The PCI Express External Cabling Specification and the SFF-8644 Specification Fixed-Side connector pinout lane mapping, although very similar, is not the same as the SAS SFF-8449 pinout lane mapping. Although the cables are compatible, attention is drawn to the fact that lanes 0 and 1 are swapped in the PCI Express definition. The SFF-8644 lane configuration facilitates simplified routing and lane ordering which is not important for formation of SAS wide ports. ~~It is however,~~ However, it is critical to PCI Express ~~in order~~ to enable lanes to be combined into x4, x8, and x16 links.



Observation: The connector used for PCI Express is also used by other common technologies; however, the pinouts are not the same. Designers should use caution when wiring the PCI Express lanes 0 and 1 to avoid inverting the lane sequence.

Table 5-15-1. Fixed-Side Connector Pinout

Row	Column								
	9	8	7	6	5	4	3	2	1
D	GND	PETn2	PETp2	GND	PETn1	PETp1	GND	MGTPWR	PWR
C	GND	PETn3	PETp3	GND	PETn0	PETp0	GND	CMISDA	CMISCL
B	GND	PERn2	PERp2	GND	PERn1	PERp1	GND	CBLPRSNT#	PWR
A	GND	PERn3	PERp3	GND	PERn0	PERp0	GND	CINT#	CADDR




Figure 5-15-1. SFF-8644 Fixed-Side Connector in Fixed-Side Cage

Table 5-25-2. Fixed-Side Multi-connector Lane

	Row	Column								
		9	8	7	6	5	4	3	2	1
Connector 0	D	GND	PETn2	PETp2	GND	PETn1	PETp1	GND	MGTPWR	PWR
	C	GND	PETn3	PETp3	GND	PETn0	PETp0	GND	CMISDA_0	CMISCL_0
	B	GND	PERn2	PERp2	GND	PERn1	PERp1	GND	CBLPRSNT#_0	PWR
	A	GND	PERn3	PERp3	GND	PERn0	PERp0	GND	CINT#_0	CADDR_0
Connector 1	D	GND	PETn6	PETp6	GND	PETn5	PETp5	GND	MGTPWR	PWR
	C	GND	PETn7	PETp7	GND	PETn4	PETp4	GND	CMISDA_1	CMISCL_1
	B	GND	PERn6	PERp6	GND	PERn5	PERp5	GND	CBLPRSNT#_1	PWR
	A	GND	PERn7	PERp7	GND	PERn4	PERp4	GND	CINT#_1	CADDR_1
Connector 2	D	GND	PETn10	PETp10	GND	PETn9	PETp9	GND	MGTPWR	PWR
	C	GND	PETn11	PETp11	GND	PETn8	PETp8	GND	CMISDA_2	CMISCL_2
	B	GND	PERn10	PERp10	GND	PERn9	PERp9	GND	CBLPRSNT#_2	PWR
	A	GND	PERn11	PERp11	GND	PERn8	PERp8	GND	CINT#_2	CADDR_2
Connector 3	D	GND	PETn14	PETp14	GND	PETn13	PETp13	GND	MGTPWR	PWR
	C	GND	PETn15	PETp15	GND	PETn12	PETp12	GND	CMISDA_3	CMISCL_3
	B	GND	PERn14	PERp14	GND	PERn13	PERp13	GND	CBLPRSNT#_3	PWR
	A	GND	PERn15	PERp15	GND	PERn12	PERp12	GND	CINT#_3	CADDR_3




Figure 5-25-2. Four Fixed-Side x4 Connector Configuration

Connectors must expand to the right of connector 0 when viewed from the external connector face, component side up, as shown in Figure 5-3.




Figure 5-35-3. Temp Image, Connector Orientation Diagram

5.4. Fixed Side Connector Labeling

This section specifies the use of optional customer-visible labeling for the Fixed-Side connector panel. It is critical to keep label size to a minimum while keeping the fields human readable.

The following attributes are used to construct a cable connector and an enclosure label:

- ❑ Signaling: The maximum signaling rate supported by the combined connectors. Allowed values are: PCIe1 (for 2.5 GT/s); PCIe2 (for 5.0 GT/s); PCIe3 (for 8.0 GT/s).
- ❑ Link Width: The maximum physical Link width of the combined connectors. Allowed values are: x1, x2, x4, x8, x12, x16.
- ❑ Port Bifurcation: (Upstream Fixed-Side only) The suffix B indicates that the combined connectors support bifurcation into smaller widths.
- ❑ Sideband Enabled: (Free-Side only) The suffix S indicates that the cable assembly supports full sideband across the cable via CML.

Examples:

- ❑ PCIe3 x16B, connector supports 8.0 GT/s up to x16, and can be bifurcated
- ❑ PCIe3 x8, connector supports 8.0 GT/s up to x8, but cannot be bifurcated
- ❑ PCIe3 x4S, assembly supports 8.0 GT/s up to x4, width, and full sideband
- ❑ PCIe3 x4, connector or assembly supports 8.0 GT/s up to x4, but assembly does not support full sideband

6. Cable Specification

6.1. SFF-8644 Cable Assembly

There are two classes of cable assemblies defined for this specification.

□ PCI Express Sideband-Isolated

The first class of cable has an isolated management interface with regard to each end of the cable assembly. In this class, there is no mechanism for external sidebands. Examples of this cable are assemblies that conform to the *SFF-8449, Shielded Cables Management Interface for SAS Specification* used in storage applications, ~~or~~ and fiber optic cables that do not provide a mechanism to effectively send CMISDA and CMISCL across the assembly. The SFF-8449 pin definitions are different than PCI Express definitions, although the cables are functionally compatible. See Section [5.3.15.2.4](#) for further details.

□ PCI Express Sideband-Enabled

The second class of cable, or PCI Express Gen3 cable, links the management interfaces at either end of the assembly together. This enables either end of the link to read the memory in the cable and for each end to communicate with the other. With the fixed ends linked via CMI, full sideband support is enabled with a higher level of configuration and other features supported.

The cable assembly must follow the pinout defined in Table 6-1. The wiring diagram in Table 6-2 indicates the internal cable connectivity from end to end. Where the connectivity does not show a direct arrow, the letters indicate which end the pin connects to at the opposite end. Power is not provided to the opposite end of the cable assembly.

Active cables must include DC-blocking capacitors between the active component and the receive pin of the connector of the cable assembly. Active cable assemblies must preserve the in-band presence detect mechanism of the SERDES as specified in the *PCI Express Base Specification*.




Figure 6-16-1. Cable Assembly with SFF-8644 Connectors

Table 6-16-1. Connector Pin assignments

Row	Column								
	9	8	7	6	5	4	3	2	1
D	GND	PETn2	PETp2	GND	PETn1	PETp1	GND	MGTPWR	PWR
C	GND	PETn3	PETp3	GND	PETn0	PETp0	GND	CMISDA	CMISCL
B	GND	PERn2	PERp2	GND	PERn1	PERp1	GND	CBLPRSNT#	PWR
A	GND	PERn3	PERp3	GND	PERn0	PERp0	GND	CINT#	CADDR

Table 6-26-2. Cable Wiring

Side 1			Side 2	
CADDR	A1	← PC ¹ x ² →	C1	CMISCL
CINT#	A2	← PC ¹ y ² →	C2	CMISDA
GND	A3	← →	C3	GND
PERp0	A4	← →	C4	PETp0
PERn0	A5	← →	C5	PETn0
GND	A6	← →	C6	GND
PERp3	A7	← →	C7	PETp3
PERn3	A8	← →	C8	PETn3
GND	A9	← →	C9	GND
PWR	B1	No Wire	D1	PWR
CBLPRSNT#	B2	No Wire	D2	MGTPWR
GND	B3	← →	D3	GND

Side 1			Side 2	
PERp1	B4	↔	D4	PETp1
PERn1	B5	↔	D5	PETn1
GND	B6	↔	D6	GND
PERp2	B7	↔	D7	PETp2
PERn2	B8	↔	D8	PETn2
GND	B9	↔	D9	GND
CMISCL	C1	↔ x ² PC ¹ →	A1	CADDR
CMISDA	C2	↔ y ² PC ¹ →	A2	CINT#
GND	C3	↔	A3	GND
PETp0	C4	↔	A4	PERp0
PETn0	C5	↔	A5	PERn0
GND	C6	↔	A6	GND
PETp3	C7	↔	A7	PERp3
PETn3	C8	↔	A8	PERn3
GND	C9	↔	A9	GND
PWR	D1	No Wire	B1	PWR
MGTPWR	D2	No Wire	B2	CBLPRSNT#
GND	D3	↔	B3	GND
PETp1	D4	↔	B4	PERp1
PETn1	D5	↔	B5	PERn1
GND	D6	↔	B6	GND
PETp2	D7	↔	B7	PERp2
PETn2	D8	↔	B8	PERn2
GND	D9	↔	B9	GND

Notes:

1. **Paddle Card (PC)** signals are only connected to the near-end backshell and paddle card. They have no wire traversing the cable assembly.
2. **x** and **y** indicate that the signals connect by wire to another location relative to the bulk wire for side-band enable cables. For cables that are not side-band enabled, these signals do not have cable wires and only connect to the near-end paddle card.



Note: There is no specific x1 cable assembly for this specification. Links smaller than x4 must use the x4 cable assembly and the link must downtrain. For ports smaller than x4, the port must begin with Lane 0.

Refer to the *SFF-8644, Mini Multilane 12 Gbs 8/4x Shielded Connector Specification* for the mechanical requirements. Electrical performance is discussed in Section 6.2. Cable assemblies must provide the information specified in the memory map described in Table 6-3 and Table 6-4.

6.2. Cable Management Interface

Cable assemblies operating at 8.0 GT/s influence the link equalization process since a predetermined equalization setting is not compatible with all cable assemblies. The cable assembly must include a memory location to provide, at a minimum, the loss information for the cable such that both the Upstream and Downstream Subsystems may derive appropriate equalization starting points for setting the Transmitter and Receiver Presets. The specifics for the memory map fields are given in Section 6.2.2.3.

The management interfaces are limited to four electrical loads and four addresses. These include a device in the Upstream Fixed-Side portion of the link, the Downstream Fixed-Side portion of the link, and each end of the cable assembly. These must be isolated from any other management interfaces in the system to limit conflicts and incompatibilities.

All cable assemblies must support CBLPRSNT#. Monolithic cable assemblies (those designed to be larger than x4) need only support one set of sideband signals, except CBLPRSNT#, to the Subsystem. The sidebands must be in the connector in which Lane 0 is housed. Manufacturers need to ensure that the loss information is consistent across the entire assembly. For cable assemblies that provide for the monitoring of the cable components, the cable assemblies need to implement the Cable Management Interface lines for each x4 connector and memory devices at both ends of the cable to provide the monitoring and status data. If memory devices are not addressable for each end of the cable assembly, the status fields for monitoring the cable plug will not be available to both ends for system monitoring.

~~For cables that exceed the capacitive loading limits given in Table 2-2, the cable assembly may need to provide mechanisms to ensure the rise time requirements are met.~~

There are circuits and components required to support functions of the cable. These are defined per sideband³; see Chapter 2.



Implementation Note: For cables that exceed the capacitive loading limits given in Table 2 2, the cable assembly may need to provide mechanisms to ensure the rise time requirements are met.

6.2.1. Memory Map

6.2.2. Memory Map Architecture

This section defines the non-volatile memory map and protocol utilized for external cable identification and management. The cables² memory device address (CADDR) is specified in Section 2.4. The LSB of the memory is configurable based on the method described there. However, cable

assemblies not supporting manageability features need only have configuration data stored at address A0h and this does not need to be configurable via CADDR.

Unless otherwise stated, all informative fields must contain accurate data. Using a value of 0 to represent a field that is unspecified is invalid. All reserved bytes and bits must be filled with logical zeroes.

For future expansion, it is recommended to keep a provision for the upper bytes 128-255 to be swapped in multiple pages, as is conventional in other memory map implementations. Which actual memory bytes are seen in addresses 128-255 then depends on the page number written in byte 127. Figure 6-2 shows the structure of the PCI Express cable CMI memory map.




Figure 6-26-2. Structure of the PCI Express Cable CMI Memory Map

The most significant bit (MSB) in a byte is bit 7, and the least significant bit (LSB) is bit 0. Unless otherwise specified, for fields that are two bytes long, the MSB is bit 7 of the higher byte number, and the LSB is bit 0 of the lower byte number.

Fields are defined as Read Only (RO), Write Only (WO), or Read/Write (R/W).

6.2.2.1. Bytes 0-127 (Lower Memory)

Table 6-3 lists the lower memory fields for Page 0.

Table 6-36-3. Page 0 Lower Memory Fields

Byte	Description	Values	R/W	Notes
0	Cable Identifier	0Fh – x4, 10h – x8, 00h – x16	RO	SFF-8636 cables use 00h as an unidentified cable type. PCI Express uses this in combination with Byte 112, bit 0:2 to indicate a x16 cable. SFF-8636 cables use 00h as an undefined cable type. PCI Express uses this in combination with Byte 112, bit 0:2 to indicate a x16 PCI Express cable.
1	Reserved			
2	Status Indicators	Bits 0 - Data_Not_Ready Bit 1 – Int# Bit 2 - Flat_mem Bits 3:7	RO	See Section 6.2.2.3.1
3	Interrupt Flags- LOS	For all bits, a value of 0 indicates de-asserted. Bit 0 – Rx0 Bit 1 – Rx1 Bit 2 – Rx2 Bit 3 – Rx3 Bit 4 – Tx0 Bit 5 – Tx1 Bit 6 – Tx2 Bit 7 – Tx3	RO	See Section 6.2.2.3.2, Optional
4-5	Reserved			
6	Interrupt Flags – Temp Alarm	Bits 0:3 – Reserved Bit 4 – Temp Low Warning Bit 5 – Temp High Warning Bit 6 – Temp Low Alarm Bit 7 – Temp High Alarm	RO	See Section 6.2.2.3.3, Optional
7	Interrupt Flags – Voltage Alarm	Bits 0:3 – Reserved Bit 4 – PWR Low Warning Bit 5 – PWR High Warning Bit 6 – PWR Low Alarm Bit 7 – PWR High Alarm	RO	See Section 6.2.2.3.4, Optional
8-21	Reserved			
22-23	Module Monitors – Temperature	Temperature in °C. Byte 22 is the upper 8 bits of 16-bit value. Byte 23 is the lower 8 bits of the 16-bit value	RO	See Section 6.2.2.3.5 Optional

Byte	Description	Values	R/W	Notes
24-25	Reserved			
26-27	Module Monitors – Supply Voltage MSB	Integer number in Hex	RO	Internally measured supply voltage, See Section 6.2.2.3.6. Optional
28-107	Reserved			
108-109	Propagation Delay		RO	See Section 6.2.2.3.7
110	Reserved			
111	Cabled PCI Express Capabilities 1	Legacy Port Adapter	RO	See Section 6.2.2.3.8
112	Cabled PCI Express Capabilities 2		RO	See Section 6.2.2.3.9
113-118	Reserved			
119-122	Password Change		WO ¹	See Section 6.2.2.3.10, Optional
123-126	Password Entry		WO ¹	See Section 6.2.2.3.10, Optional
127	Page Select		R/W	See Section 6.2.2.3.11
¹ Or Reserved				

6.2.2.2. Bytes 128-255 (“Page 0, Upper Memory”)

Table 6-4 lists the Upper Memory Fields for Page 0.

Table 6-4-4. Page 0 Upper Memory Fields

Byte	Description	Values	R/W	Notes
128	Cable Identifier	Must match Byte 0	RO	See Byte 0 in Table A-3
129-146	Reserved		RO	
147	Cable Technology	Bits 0:3 Reserved Bits 4:7 Cable Technology Type	RO	Cable equalization and medium See Section 6.2.2.3.12
148-163	Vendor Name	ASCII string (16 characters)	RO	See Section 6.2.2.3.13
164	Reserved			
165-166	PCI-SIG Vendor ID	Hex number (two bytes)	RO	See Section 6.2.2.3.14

Byte	Description	Values	R/W	Notes
167	Reserved			
168-183	Vendor Part Number	ASCII string (16 characters)	RO	See Section 6.2.2.3.15
184-185	Vendor Revision	ASCII string (2 characters)	RO	See Section 6.2.2.3.16
186	Attenuation 1.25 GHz	Hex number	RO	See Section 6.2.2.3.17
187	Attenuation 2.5 GHz	Hex number	RO	See Section 6.2.2.3.17
188	Attenuation 4.0 GHz	Hex number	RO	See Section 6.2.2.3.17
189	Attenuation 8.0 GHz	Hex number	RO	See Section 6.2.2.3.17
190	Maximum case temperature	Hex number	RO	See Section 6.2.2.3.18
191	Checksum - Base	Checksum on Bytes 128-190	RO	See Section 6.2.2.3.19
192-194	Reserved			
195	Options	Bits 0:7 – Reserved	RO	See Section 6.2.2.3.20
196-211	Vendor Serial Number	ASCII string (16 characters)	RO	See Section 6.2.2.3.21
212-217	Vendor Date Code	ASCII string (yymmdd)	RO	See Section 6.2.2.3.22
218-219	Vendor Lot Code	Hex number (two bytes)	RO	See Section 6.2.2.3.23
220-222	Reserved	Hex number (two bytes)		
223	Checksum - Extended	Checksum for Bytes 192-222	RO	See Section 6.2.2.3.24
224-255	Vendor Specific Information	Vendor Specific (32 Bytes)	RO	See Section 6.2.2.3.25

6.2.2.3. Field Descriptions

This section defines the non-volatile memory map and protocol utilized for external cable identification and management.

~~All fields contain non-negative values. All fields are specified as a positive value.~~

6.2.2.3.1. Bytes 1-2, Status Indicators

The Status Indicator~~orsion~~ field indicates the status of cable management information. Table 6-5 lists the bit descriptions.

Table 6-56-5. Status Indicators

Bit	Description
0	Data_Not_Ready: Asserted until a valid set of monitor readings
1	Int# (Optional): Digital indicator reflecting CINT#
2	Flat_mem
3:7	Reserved

6.2.2.3.2. Byte 3, Interrupt Flags – LOS

The Byte 3 latched bits indicate a loss of signal per lane.

The asserted value must be 1b. Once a bit is Set, CINT# is asserted and the value is retained until a read operation of the bit or the free side is reset by removal and assertion of MGTPWR. The bit is permitted to be set again any time the criteria for the assertion of the signal are met.

- Bits 0 to 3 indicate the Rx lanes, with 0 being Lane 0 and the bit location incrementing per lane.
- Bits 4 to 7 indicate the Tx lanes, with bit 4 being Lane 0 and the bit location incrementing per lane.

6.2.2.3.3. Byte 6, Interrupt Flags – Temp Alarm

The Byte 6 latched bits indicate a temperature condition.

The asserted value must be 1b. Once a bit is Set, CINT# is asserted and the value is retained until a read operation of the bit or the Free-Side is reset by removal and assertion of MGTPWR. The bit is permitted to be set again any time the criteria for the assertion of the signal are met.

- Bits 0-3 Reserved
- Bits 4 and 5 provide warnings and indicate low and high conditions respectively.
- Bits 6 and 7 provide alarms and indicate low and high conditions respectively.

6.2.2.3.4. Byte 7, Interrupt Flags – Voltage Alarm

The Byte 7 latched bits indicate a voltage condition.

The asserted value must be 1b. Once a bit is Set, CINT# is asserted and the value is retained until a read operation of the bit or the Free-Side is reset by removal and assertion of MGTPWR. The bit is permitted to be set again any time the criteria for the assertion of the signal are met.

□ Bits 0-3 Reserved

□ Bits 4 and 5 provide warnings and indicate low and high conditions respectively.

□ Bits 6 and 7 provide alarms and indicate low and high conditions respectively.

6.2.2.3.5. Bytes 22-23, Module Monitors – Temperature

Internally measured Free-Side device temperatures are represented as a 16-bit signed twos' complement value in increments of 1/256 degrees Celsius, yielding a total range of -128°C to +128°C that is considered valid between -40°C and +125°C. Temperature accuracy is Vendor-Specific but must be better than $\pm 3^\circ\text{C}$ over specified operating temperature and voltage. Placement of the temperature sensor is vendor-specific, however the temperature reported in Bytes 22-23 should correspond to the estimated case temperature.

6.2.2.3.6. Bytes 26-27, Module Monitors – Supply Voltage

Internally measured Free-Side device supply voltages are represented as a 16-bit unsigned integer with the voltage defined as the full 16-bit value (0 to 65535) with LSB equal to 100 μV , yielding a total measurement range of 0 to +6.55 V. Practical considerations, to be defined by Free-Side device manufacturer, tends to limit the actual bounds of the supply voltage measurement. Accuracy is vendor-specific, but must be better than $\pm 3\%$ of the manufacturer's nominal value over specified operating temperature and voltage.

6.2.2.3.7. Bytes 108-109, Propagation Delay

These two bytes indicate the range of the one-way propagation delays through the cable assembly of all the signal pairs. The propagation delays all fall between a lower bound and an upper bound. The lower bound is defined as the largest multiple of 250 ps that is less than or equal to the propagation delay of the fastest signal pair in the cable. In other words, it is the smallest propagation delay, rounded down using a floor function to a multiple of 250 ps. The upper bound is defined as the smallest multiple of 250 ps that is greater than or equal to the propagation delay of the slowest signal pair in the cable. In other words, it is the largest propagation delay, using a ceiling function rounded up to a multiple of 250 ps.

The upper and lower bounds are expressed by a combination of two parameters; Base and Span, in the Propagation Delay field, as shown in Figure 6-3.

Byte	108								109							
Bit	7 (msb)	6	5 (lsb)	4 (msb)	3	2	1	0	7	6	5	4	3	2	1	0 (lsb)
Contents	Span								Base							

Figure 6-36-3. Bytes 108-109, Propagation Delay

The following equations define the parameters:

- $\text{Base} = (\text{Lower Bound}) / (250 \text{ ps})$
- $\text{Span} = 0$, for $(\text{Upper Bound} - \text{Lower Bound}) \leq 250 \text{ ps}$
- $\text{Span} = (\text{Upper Bound} - \text{Lower Bound}) / (250 \text{ ps}) - 2$, for $(\text{Upper Bound} - \text{Lower Bound}) > 250 \text{ ps}$

The Lower Bound, Upper Bound, Base, and Span values for some example cables are shown in Table 6-6.

Table 6-6-6. Lower Bound, Upper Bound, Base, and Span Values

Parameter	Cable 1	Cable 2	Cable 3	Cable 4
Fastest Pair's Propagation Delay	10.2 ns	133.6 ns	794.8 ns	2047.9 ns
Slowest Pair's Propagation Delay	12.1 ns	133.7 ns	795.2 ns	2049.2 ns
Lower Bound	10.00 ns	133.50 ns	794.75 ns	2047.75 ns
Upper Bound	12.25 ns	133.75 ns	795.25 ns	2049.25 ns
Upper Bound - Lower Bound	2.25 ns	0.25 ns	0.5 ns	1.5 ns
Base	40 (28h)	534 (216h)	3179 (C6Bh)	8191 (1FFFh)
Span	7 (111b)	0 (000b)	0 (000b)	4 (100b)

A cable management controller computes the Upper and Lower Bound for a cable based on the following equations:

- $\text{Lower Bound} = \text{Base} * 250 \text{ ps}$
- $\text{Upper Bound} = (\text{Base} + \text{Span} + 2) * 250 \text{ ps}$



Observation: The lowest Upper Bound that is expressed by the Span parameter is 500 ps more than the Lower Bound.

Observation: Since actual propagation delays do not fall on exact 250 ps boundaries, virtually all cables report a range of propagation delays that is up to 250 ps longer than the actual cable skew. Thus, a cable that has almost 2 ns of skew is most likely to report a 2.25 ns range around its fastest and slowest propagation delays. Due to this, a cable management controller should not compare a cable's own reported Upper Bound and Lower Bound with each other when determining cable compatibility for port aggregation. It should compare other cables reported Upper Bound with every other cable reported Lower Bound. Any difference greater than +2 ns indicates that the combination has excessive skew and should not be aggregated.

6.2.2.3.8. Byte 111, Cabled PCI Express Capabilities 1

Table 6-7 lists the bit descriptions for the Cabled PCI Express Capabilities 1.

Table 6-7. Cabled PCI Express Capabilities 1

Bit	Description
0	2.5 GT/s, must be set to 1b
1	1b if 5.0 GT/s capable, else 0b
2	1b if 8.0 GT/s capable, else 0b
3-6	Reserved
7	Legacy Port Adapter

6.2.2.3.9. Byte 112, Cabled PCI Express Capabilities 2

Table 6-8 lists the bit descriptions for the Cabled PCI Express Capabilities 2

Table 6-8. Cabled PCI Express Capabilities 2

Bit	Value	Description
0-2	000b	x1
	001b	x2
	010b	x4
	011b	x8
	100b	x12
	101b	x16
		All others Reserved
3-6		Reserved
7		Support for Sideband Messages

6.2.2.3.10. Bytes 119-126 Password Entry and Change

Bytes 119-126 are used for the optional password entry function. Password entry bytes are write only and are retained until power down, reset, or rewritten by the Fixed-Side. If this function is not supported, all these bytes are Reserved. Byte 123 bit 7 is the most significant bit and Byte 126 bit 0 is the least significant bit. Similarly, Byte 119 bit 7 is the most significant bit and Byte 122 bit 0 is the least significant bit.

Password Entry field is field retained until power down, reset, or rewritten by the Fixed-Side.

The Password Change field is discarded once the written value has updated the password. These fields are write-only.

Cable vendors are permitted to use this function to implement write protection of identification fields and other read only information. Password access must not be required to access Free-Side device data in either the lower memory page 00h or upper page 00h. Note that multiple manufacturer passwords are permitted to be defined to allow selective access to read or write to various sections of memory, as allowed above.

Fixed-Side Subsystem manufacturer and cable manufacturer passwords must be distinguished by the high order bit (bit 7, byte 123). All Fixed-Side Subsystem manufacturer passwords must fall in the range of 0000 0000h to 7FFF FFFFh, and all cable manufacturer passwords in the range of 8000 0000h to FFFF FFFFh. Fixed-Side Subsystem manufacturer passwords must be initially set to 0000 1011h in new cables.

Fixed-Side system manufacturer passwords are permitted to be changed by writing a new password in Bytes 119-122, when the correct current Fixed-Side manufacture password has been entered in 123-126; with the high order bit being ignored and forced to a value of 0 in the new password. When the Password Change field is written, the new password value must immediately come into effect, blocking all password protected information, until the Password Entry field contains the new password value.

The current password value(s) must be retained through power cycling or reset. It is strongly recommended that cable assembly vendors provide some vendor-specific method to return the password value(s) to its original factory value.

The Password Entry field must be set to 0000 0000h on power up and reset.

6.2.2.3.11. Byte 127 Page Select

Byte 127 is used to select the upper page. A value of 00h indicates upper memory page 00h is available to be mapped to locations 128 to 255. All other values are reserved for future use. Page 02 is optionally provided as user writable EEPROM. The Fixed-Side reads or writes this memory for any purpose. If Page 00h Byte 129 bit 4 is set, the first 10 bytes of Page 02h Bytes 128-137 are used to store the Common Language Equipment Identification (CLEI) code for the free side device.

[The](#) Common Language Equipment Identifier is SFF-8636.

6.2.2.3.12. Byte 147, Cable Technology

The Cable Technology register indicates to the Fixed-Side Subsystem what type of cable is inserted. This information is used in conjunction with Bytes 186-189 to determine optimal equalizer settings for transmitters and receivers. FSS designers should see Section 4.2.1.1 to understand how to interpret that data in Table 6-9 and use it for configuration of the PCI Express devices.

Table 6-96-9. Cable Technology Type

Value	Description
0h	Active Optical Cable
8h	Undefined/Other
Ah	Copper Cable, Unequalized
Bh	Copper Cable, Passively Equalized
Ch	Copper Cable, TX and RX Limiting Active Equalizers
Dh	Copper Cable, Rx Limiting Active Equalizer
Fh	Copper Cable, Active Linear Equalizer
	All other encodings are Reserved

6.2.2.3.13. Bytes 148-163, Vendor Name

The Vendor Name is a 16-byte field that contains ASCII characters, left-aligned, and padded on the right with ASCII spaces (20h). The Vendor Name must be the full name of the corporation, a commonly accepted abbreviation of the name of the corporation, or the stock exchange code for the corporation. A value of all zeroes in the field indicates that the Vendor Name is unspecified. At least one of the Vendor Name or the PCI-SIG Vendor ID must contain valid non-zero data.

6.2.2.3.14. Bytes 165-166, PCI-SIG Vendor ID

This is the 2-byte Vendor ID, allocated by the PCI-SIG. A value of all zeroes indicates that the Vendor ID is unspecified.

6.2.2.3.15. Bytes 168-183, Vendor Part Number

The Vendor Part Number is a 16-byte field that contains ASCII characters, left aligned and padded on the right with ASCII spaces (20h). This field defines the vendor part number or product name. A value of all zeroes in the field indicates that the Vendor Part Number is unspecified.

6.2.2.3.16. Bytes 184-185, Vendor Revision

The Vendor Revision is a 2-byte field that contains ASCII characters, left aligned, and padded on the right with ASCII spaces (20h) and defines the vendor's product revision number. A value of all zeroes in the field indicates that the Vendor Revision is unspecified.

6.2.2.3.17. Bytes 186-189, Attenuation at Frequency

The data in these fields indicate the loss of the cable assembly and mated connectors in increments of 1 dB for each frequency. Test fixture loss must not be included in this data. A value of all zeroes indicates that the attenuation is not known or is unavailable. It is recommended that values of all zeros be used for cables in which the cable isolates the electrical waveform from either of the Fixed-Sides. Examples of this include active optical cables (Cable Technology Type field value of 0h) and cable with limiting amplifiers at both the Tx and Rx ends of the cable (Cable Technology Type field value of Ch).

6.2.2.3.18. Byte 190, Maximum Case Temperature

Max case temperature in degrees Celsius. The value in degrees is translated directly to a hex value. For an unspecified value (e.g., 00h), a limit of 70°C is assumed. A value representing 0°C is not allowed.



Note: An unspecified value is treated as 00h and this is interpreted as 70°C rather than 0°C.

6.2.2.3.19. Byte 191, Checksum-Base

A check code to verify the first 63 bytes of the upper memory page. The value is the low order eight bits of the sum of the contents of bytes 128-190.

6.2.2.3.20. Byte 195, Options

This field is reserved for future use to indicate cable options.

6.2.2.3.21. Bytes 196-211, Vendor Serial Number

The Vendor Serial Number is a 16-byte field that contains ASCII characters, left aligned, and padded on the right with ASCII spaces (20h) and defines the vendor's serial number for the Free-Side device. A value of all zeroes in the field indicates that the Vendor Serial Number is unspecified.

6.2.2.3.22. Bytes 212-217, Vendor Date Code

A six-byte field containing the vendor's date code in ASCII. Bytes 212-213 contain the two-year digits (00=2000, 99=2099). Bytes 214-215 contain the month (01=Jan, 12=Dec). Bytes 216-217 contain the day of the month (01-31).

6.2.2.3.23. Bytes 218-219, Vendor Lot Code

Vendor defined lot code. This is permitted to be all zeros.

6.2.2.3.24. Byte 223, Checksum-Extended

A check code to verify the 31 bytes above the Checksum-Base in the upper memory page. The value is the low order 8 bits of the sum of the contents of bytes 192-222.

6.2.2.3.25. Bytes 224-225, Vendor Specific Information

Vendor specified data.

6.2.3. Non-Volatile Memory Specification

The memory transaction timings are given in Table 6-10.

Table 6-10. Non-volatile Memory Specifications

Parameter	Symbol	Min	Max	Unit	Conditions
Complete Single or Sequential Write	t_{WP}		40	ms	Complete (up to) 4-byte Write
Serial Interface Clock Hold off (Clock Stretching)	$t_{SCL_HOLDOFF}$		$T_{CMI_CLK_HOLD}$ (see Table 2-2)	μs	Maximum time the cable module is permitted to hold the CMISCL line low before continuing with a read or write operation
Endurance (Write Cycles)		50,000		Cycles	70° C

6.2.4.6.3. Cable Electrical Performance



Note: Values and limits for frequency domain parameters are typically listed in this document as Loss Values and are therefore positive. However, since they are losses, the actual measured values, either simulated or measured, would be negative. Table 6-11 lists the electrical performance requirements for PCI Express external cables.

Cable assembly compliance testing is done using frequency domain techniques. The insertion loss, return loss, and crosstalk requirements are provided starting in Section 6.3.1.6.2.4.1. Cable assembly time domain testing is performed against the eye diagrams in lieu of the frequency domain parameters. This allows trade-offs between the cable assembly interconnect parameters and implementation variables, such as equalization. Table 6-11 lists a summary of the cable assembly differential characteristics and is included for convenience. The referenced section contains the specification requirement.

Table 6-11. Cable Assembly Differential Characteristics Summary

Description	Values for 8.0 GT/s ¹	Unit	Reference
Maximum Insertion Loss at 6.5 GHz	11.4	dB	See Section 6.2.4.1
Maximum Insertion Loss at 4.5 GHz	8.4		
Maximum Insertion Loss at 8.0 GHz	14.6		
Maximum Insertion Loss at 4.0 GHz	9.67.8		
Maximum Insertion Loss at 2.5 GHz	7.35.9		
Maximum Insertion Loss at 1.25 GHz	4.94.0		
Maximum Insertion Loss at 625 MHz	3.32.7		
Minimum Return Loss at 8.0 GHz	TBD	dB	See Section 6.2.4.2
Minimum Return Loss at 4.0 GHz	TBD6.5		
Minimum Return Loss at 2.5 GHz	TBD12		
Minimum Return Loss at 1.25 GHz	TBD12		

Description	Values for 8.0 GT/s ¹	Unit	Reference
Minimum DCMC below 2.58 GHz	TBD 18		See Section 6.2.4.9
Notes: Applies to copper cable assemblies without any form of non-linear equalization.			

6.2.4.1.6.3.1.1. Maximum Insertion Loss

6.2.4.1.1.6.3.1.1.1. 8.0 GT/s

The insertion loss for each pair in a cable assembly operating at 8.0 GT/s must meet the values using Equation 6-1.

Equation 6-16-1. 8.0 GT/s Maximum Insertion Loss

$$InsertionLoss(f) \leq 0.1600 \sqrt{f} + 0.000437 \times f \text{ dB } 100 \leq f \leq 804500$$

$$InsertionLoss(f) \leq 1.51e - 3 * f + 1.56 \text{ dB } 4500 \text{ MHz} \leq f \leq 6500$$

where f is the frequency in MHz. ~~See Equation 6-2.~~

Controlling the ISI jitter component through additional equalization within the cable assembly is allowed. Details of such implementations are beyond the scope of this specification. Figure 6-4 shows a graph of cable assembly insertion loss.





Figure 6-46-4. Cable Assembly Insertion Loss

6.2.4.2.6.3.1.2. Minimum Return Loss

6.2.4.2.1.6.3.1.2.1. 8.0 GT/s

The return loss of each pair of the cable assembly must meet the values determined using Equation 6-2. ~~See Figure 6-5 for an illustration of 8 GT/s insertion loss limit.~~

Equation 6-26-2. 8.0 GT/s Minimum Return Loss

$$\text{ReturnLoss}(f) \geq 12 \text{ dB} \quad 100 \leq f(\text{MHz}) \leq 4000$$

$$\text{ReturnLoss}(f) \geq 6.5 \text{ dB} \quad 4000 < f(\text{MHz}) \leq 8000$$

where f is the frequency in MHz.

6.2.4.3.6.3.1.3. Integrated Crosstalk Noise (ICN) Multiple Disturber NEXT Loss

~~The NEXT loss that is coupled into a receive lane is from multiple transmit lanes at the near end. Integrated Crosstalk Noise, or ICN, is a figure of merit for crosstalk. This performance metric combines near-end and far-end crosstalk, across a specific frequency range, into a single value to be compared to insertion loss.~~

~~To calculate ICN, MDNEXT and MDFEXT from the cable assembly crosstalk components that have N uniformly spaced frequencies, f_n , with a maximum frequency span, 10 MHz, over the frequency range specified in Table 6-12. Each of these is then multiplied by transmitter and receiver bandwidth filters. Each data set is then integrated and combined to calculate an RMS value for ICN for the cable assembly.~~

Table 6-12. Calculate ICN, MDNEXT and MDFEXT

Symbol	Description	Value	Units
f_b	Baud Rate	8	Gbd
A_{NE}	Amplitude of NEXT Aggressor	0.8	Volts
A_{FE}	Amplitude of FEXT Aggressor	0.8	Volts
T_{Rise_NE}	Rise & Fall Time of Near End Aggressor	46	ps
T_{Rise_FE}	Rise & Fall Time of Near End Aggressor	46	ps
F_{IL_ICN}	Frequency of Insertion Loss data used for ICN comparison	4	GHz
f_r	Receiver Bandwidth	5.82	GHz
$F_{ICN \text{ Lower Limit}}$	Lower frequency limit for ICN	50	MHz
$F_{ICN \text{ Upper Limit}}$	Upper frequency limit for ICN	6500	MHz

6.3.1.3.1. Multiple Disturber Near-End Crosstalk (MDNEXT)

MDNEXT is determined by summing the adjacent individual pair-to-pair differential NEXT loss values using Equation 6-3.

Equation 6-3. MDNEXT

$$MDNEXT(f) = \sum_{i=1}^n 10^{-NEXT(f)_i/10}$$

where f is the frequency in MHz.

$NEXT(f)_i$ is the magnitude in dB of the NEXT loss for frequency f for each data point i , and n is the number of data points.

6.3.1.3.2. Multiple Disturber Far-End Crosstalk (MDFEXT)

MDFEXT is determined by summing the adjacent individual pair-to-pair differential FEXT loss values using Equation 6-4.

Equation 6-4. MDFEXT

$$MDFEXT(f) = \sum_{i=1}^n 10^{-FEXT(f)_i/10}$$

where f is the frequency in MHz.

$FEXT(f)_i$ is the magnitude in dB of the FEXT loss for frequency f for each data point i , and n is the number of data points.

6.3.1.3.3. Weighting Functions

The weighting functions are used to act as filter functions for transmitters and receivers when combined with the MDNEXT and MDFEXT functions. The weighting functions are defined by Equation 6-5 and Equation 6-6.

Equation 6-5. Weighting Function for MDNEXT

$$W_{NE}(f_n) = \left(\frac{A_{NE}^2}{f_b}\right) \times \text{sinc}^2\left(f_n/f_b\right) \left[\frac{1}{1+(T_{Rise_NE} \cdot f_n/0.2365)^4}\right] \left[\frac{1}{1+(f_n/f_r)^8}\right] -$$

Equation 6-6. Weighting Function for MDFEXT

$$W_{FE}(f_n) = \left(\frac{A_{FE}^2}{f_b}\right) \times \text{sinc}^2\left(f_n/f_b\right) \left[\frac{1}{1+(T_{Rise_FE} \cdot f_n/0.2365)^4}\right] \left[\frac{1}{1+(f_n/f_r)^8}\right] -$$

where f_n is the frequency in GHz and f_b, f_r, T_{Rise_NE} and T_{Rise_FE} are given in Table 6-12.

The SINC function is defined by Equation 6-7.

Equation 6-7. SINC Function

$$\text{sinc}(x) = \frac{\sin(\pi x)}{(\pi x)}$$

6.3.1.3.4. Integrated Crosstalk

The near-end and far-end integrated crosstalk components are calculated by Equation 6-8 and Equation 6-9.

Equation 6-8. Near-end Crosstalk

$$\sigma_{NE} = \sqrt{2 \times \Delta f \times \sum_n W_n(f_n) \times MDNEXT(f_n)}$$

Equation 6-9. Far-end Crosstalk

$$\sigma_{FE} = \sqrt{2 \times \Delta f \times \sum_n W_n(f_n) \times MDFEXT(f_n)}$$

The total ICN for the cable assembly is calculated using Equation 6-10.

Equation 6-10. Total ICN for Cable Assembly

$$ICN = \sqrt{\sigma_{NE}^2 + \sigma_{FE}^2}$$

6.3.1.3.5. ICN Limit

The ICN is compared to the cable assembly's insertion loss at half the baud rate. Equation 6-11 determines the limits for ICN for a given cable assembly when compared to the half baud rate in dB. See Figure 6-5.

Equation 6-11

$$ICN \leq 46.57 + 5.71 \times IL(F_{IL_ICN}) \text{ mV}$$




Figure 6-5. Insertion Loss at F_{IL_ICN}

1623

~~Equation 6-4. 8.0 GT/s MDNEXT Loss Definition~~

1624 ~~$MDNextLoss(f)$ is defined as $10 \times \log \left(\sum_{i=1}^n 10^{-NextLoss(f)_i / 10} \right) dB$~~

1625 ~~where f is the frequency in MHz.~~

1626 ~~$NextLoss(f)_i$ is the magnitude in dB of the NEXT loss for frequency f for each measurement i ,~~

1627 ~~where n is the number of measurements. If the pair-to-pair NEXT loss measured between 100 MHz and~~

1628 ~~2500 MHz is greater than 50 dB over the entire frequency range, that measurement sweep does not need~~

1629 ~~to be included in the MDNEXT summation.~~

1630

~~6.2.4.4. MDFEXT~~

1631 ~~Multiple-Disturber far-end crosstalk (MDFEXT) is used to calculate the sum of all far-end crosstalk~~
1632 ~~caused by multiple aggressor lanes and originating at the far end of the channel. FEXT(f) $_i$ is the~~
1633 ~~magnitude in dB of the FEXT at frequency f for each measurement i , and n is the number of~~
1634 ~~measurements.~~

1635 ~~Equation 6-8. Determining MDFEXT~~

1636 ~~$MDFEXT(f)$ is defined as $-10 \times \log \sum_{f=1}^n 10^{-FEXT(f)/10} dB$~~

1637

~~6.2.4.5. Insertion Loss to Crosstalk Ratio~~

1638 ~~Insertion Loss to Crosstalk Ratio (ICR) is the ratio of the insertion loss to the total crosstalk.~~

1639 ~~Equation 6-9. Insertion Loss to Crosstalk~~

$$ICR(f) = 10 * LOG_{10} \left(10^{\Lambda \left(\frac{MDNEXT(f)}{10} \right)} + 10^{\Lambda \left(\frac{MDFEXT(f)}{10} \right)} \right) - IL \text{ dB } 100 < f < 8000$$

where f is the frequency in MHz.

6.2.4.6.6.3.1.4. Maximum Pair-to-Pair Skew

The difference in propagation delay, or skew, between all cable assembly pair combinations, must not exceed the cable skew limit, S_C in Section 3.2.4.

6.2.4.7.6.3.1.5. Maximum Intra-Pair Skew

6.2.4.7.1.6.3.1.5.1. 8.0 GT/s

There is no explicit intra-pair skew requirement for 8.0 GT/s.

6.2.4.8.6.3.1.6. Maximum DC Resistance for Auxiliary Signals

The resistance of any conductor used for Auxiliary signals must not exceed 10 Ω .

6.2.4.9.6.3.1.7. Differential-to-Common Conversion Loss

The minimum differential-to-common mode signal conversion loss in the cable assembly must be determined using [Equation 6-12](#)~~Equation 6-8~~. This specification parameter only applies to cable assemblies operating at 8.0 GT/s.

Equation 6-12~~6-10~~. 8.0 GT/s Differential-to-Common Conversion Loss

$$DCMC(f) \geq 18dB \quad 100 \leq f \leq 8000$$

where f is the frequency in MHz

6.2.4.10.6.3.1.8. CMI Line Capacitance

The total capacitance for a CMI line within the completed cable assembly is defined by $C_{CMI_CABLE_LOAD}$. If the cable assembly bus capacitance exceeds $C_{CMI_CABLE_LOAD}$, active devices are to be employed, but all the AC specifications must still be met. For impedance targets of the CMI lines, see Table 2-2. The logic threshold requirements that must be met by active devices in the cable are in Table 2-1.

6.2.4.11.6.3.1.9. HiPot Requirement

The minimum HiPot requirement for cable assemblies must be 240 VDC for 100 ms.

6.3.6.3.1.10. Shielding

All cable assemblies must provide 360 degrees shielding from end-to-end. EMI gaskets are required at the mating interface to reduce EMI emissions to an acceptable level.

6.4. Cable Assembly Labeling

This section specifies the customer-visible labeling of the cable assembly. It is critical to keep label size to a minimum while keeping the fields human readable.

The following attributes are used to construct a cable assembly:

- ❑ PCI Express latch pull tab to be Pantone 354U Green.
- ❑ Signaling: The maximum signaling rate supported by the combined connectors. Allowed values are: PCIe1 (for 2.5 GT/s); PCIe2 (for 5.0 GT/s); PCIe3 (for 8.0 GT/s).
- ❑ Link Width: The maximum physical Link width of the cable assembly. Allowed values are: x1, x2, x4, x8, x12, x16.
- ❑ Sideband Enabled: (Free-Side only) The suffix S indicates that the cable assembly supports full sideband across the cable via CMI.

Examples:

- ❑ PCIe3 x8S, assembly supports 8.0 GT/s up to x8, width, and full sideband
- ❑ PCIe3 x4, connector or assembly supports 8.0 GT/s up to x4, but assembly does not support full sideband

1684

1685 **7. Other Considerations**

1686 When expanding a system using a cabled interface, there are platform considerations that system
1687 vendors and users need to be aware of. While individual links function, the system must deal with
1688 configurations that may be beyond the scope of perceived usage models. Items such as bus
1689 enumeration, memory allocations, and platform timing need to be understood for the entire system
1690 to work correctly. Some of these considerations are presented here for awareness.

1691 **7.1. Discovery of PCI Express Topology**

1692 A cabled interface to an existing system allows for a sizable increase in the number of system
1693 components. In the case of a cabled PCI Express interface, the number of busses grow to approach
1694 the limit of the specification. Some system firmware are not designed to handle this seamlessly and
1695 could be a problem for the host system. System integrators and end users should be aware that it is
1696 possible to exceed what the system's original firmware was designed to operate with, which leads to
1697 unknown system behaviors.

1698 **7.2. Memory Allocation**

1699 Adding additional devices to an existing system allows for a sizable increase in the number of system
1700 components, and therefore an increase in the system resources required to enable these components
1701 to operate. A key component of this is system memory. PCI devices are allocated memory upon
1702 discovery, and later they may request more. A cabled system allows for an extensive number of
1703 devices to be added to the original system and the system hardware and firmware may not be able to
1704 predictably operate with demands larger than what the original system was designed for.

7.3. Completion Timeouts of PCI Express Transactions

Adding additional devices to an existing system allows for a sizable increase in the number of system components. These components are added in a cascaded fashion through a series of PCI Express switches. There are aspects to the protocol that require operations to complete in an expected timeframe. While the standard does allow for large delays or programmable delays in many implementations, these flexibilities are not required by the standard. This may be a problem for the host system or the Downstream Subsystem. System integrators and end users should be aware that exceeding what the system's original firmware was designed to operate with leads to unknown system behaviors.

7.4. Bus Performance and Throughput for Interfaces with Long Flight Times

An active cable, either copper or fiber, allows for a significant propagation delay of a signal between devices, well beyond what the *PCI Express Base Specification* originally allowed. While the protocol deals with long delays, there are impacts to system performance. The protocol includes flow control credits to limit the amount of data in flight so as to not overflow a device's link receiver buffers. Some data transactions require an acknowledgement that data was received before the flow control credits are released so that more data transactions are sent. Independent of device performance, the delay induced by the physical transmission medium is at least twice the propagation delay (once for the data transmission and again for the return acknowledgement) before the next data transaction is sent. The protocol has some built-in checks to make sure a data transaction is not lost after it is sent. In addition to the Update FC being returned for flow control, the ACK must be returned on a finite schedule. If the data is not received within a certain window, the data transaction is sent again, or replayed. The protocol allows for a data transaction to be replayed three times, for a total of four attempts at data transmission and acknowledgement. After the fourth attempt, a correctable error occurs, and the link attempts to recover itself.

7.5. Interoperability with SAS cables

PCI Express supports the use off-the-shelf cable assemblies that are compliant with the SFF-8449 and the SFF-8636 specifications. However, the link setup information contained in the cable assembly memory device is different from that of a PCI Express specific cable assembly.

An SFF-8449 cable assembly has a $T_{\text{MGMT_INTF}}$ time of 2000 ms.

~~7.5.1.1.1.~~



Note: SFF-8449 pin definitions are different than PCI Express definitions, though the cables are functionally compatible. See Section ~~5.3.15.2.4~~ for further details.

~~7.5.1.1.2.~~

The CMI controller in the Fixed-Subsystems must translate the SFF-8636 data into a format that is useful for programming the PCI Express device. The controller determines if the cable is PCI Express by checking the Identifier field (Byte 0).

It should be noted that SFF-8449 cables do not support sideband messages and any behaviors that depend on sideband messages (e.g. boot sequencing, local resets, wake) need to be adjusted, when this type of cable is used to create the cabled link.

Key differences in the memory map relating to configuration data are:

□ Bytes 108-109

In SFF-8636, initially developed for SAS, Bytes 108-109 contain propagation delay in multiples of 10 ns. However, in a PCI Express cable, Bytes 108-109 contain propagation delay and lane-to-lane cable skew in multiples of 250 ps, as defined in Section 6.2.2.3.7.

□ Bytes 111-112

In SFF-8644, Bytes are all zeroes, however, PCI Express cables use them to describe PCI Express capabilities.

□ Byte 147

PCI Express uses a subset of the options defined in an SFF-8449 cable. Values of 8h, Ah, Bh, Ch, Dh, and Fh are the same for both types. In an SFF-8449 cable, a value of 0h specifies a particular type of optical driver and values of 1h-7h and 9h all indicate other types of optical drivers. PCI Express is not required to distinguish between types of optical drivers and only the value of 0h indicates an optical driver is used.

□ Bytes 165-166

These will not contain a the PCI-SIG ID Vendor ID.

□ Byte 186-189

Indicate performance characteristics of the cable. Detailed recommendations to implement PCI Express devices settings are beyond the scope of this specification. However, differences are highlighted to allow firmware designers to implement solutions if as they choose to.

• Byte 186

Specify either wavelength or copper attenuation. Copper attenuation specified at 0 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, is 1b. Byte 147 indicates if this is a copper or optical interface. If it is optical, this field is not relevant for configuration. If it is copper, this is the bit rate, not the Nyquist frequency, and is likely to indicate more loss than what is present at Nyquist.

• Byte 187

Specify either wavelength or copper attenuation. Copper attenuation specified at 0 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, is 1b. Byte 147 indicates if this is a copper or optical interface. If it is optical, this field is not relevant for configuration. If it is copper, this is the bit rate, not the Nyquist frequency, and is likely to indicate more loss than what is present at Nyquist.

• Byte 188

Specify either wavelength or copper attenuation. Copper attenuation specified at 7 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, is 1b. Byte 147 indicates if this is a copper or optical interface. If it is optical, this field is not relevant for configuration. If it is copper, this frequency is a 12% below the bit rate, or 75% above the Nyquist frequency.

1781 • Byte 189
1782 Specify either wavelength or copper attenuation. Copper attenuation specified at
1783 12 GHz if Byte 147 is Ah-Fh and Byte 133, bit 6, is 1b. If it is copper, this is the bit rate, not
1784 the Nyquist frequency, and is likely to indicate more loss than what is present at Nyquist. If it
1785 is copper, this frequency is a 25% below the bit rate, or 50% above the Nyquist frequency.

~~7.6. PCIe Registers~~

1786 —
1787 —

1788 —
1789 —

1790

1791 **APPENDIX A. Acknowledgements**

1792 The following persons were instrumental in the development of the *PCI Express External Cabling*
1793 *Specification* (Company affiliation listed is at the time of specification contributions):

1794

1795